Author: Amy: Refs. are not cited in numerical order in this chapter.

chapter twelve

Adaptive routing and wavelength assignment in all-optical networks: the role of wavelength conversion and virtual circuit deflection

Author: Please verify that the terms "full-wavelength conversion" and "no-wavelength conversion" should be hyphenated throughout this chapter.

Emmanuel A. Varvarigos
University of Patras
Theodora Varvarigou
National Technical University of Athens
Evangelos Verentziotis
National Technical University of Athens

Contents

0-8493-XXXX-X/02/$0.00+$1.50
© 2003 by CRC Press LLC

12.1 Introduction

Session establishment in an all-optical network involves two kinds of decisions:

1. The selection of a route, or sequence of hops, that the session must traverse
2. For each hop along the route and according to the wavelength conversion capability of the corresponding switching node, the selection of a wavelength on which the session will be carried for that hop

This chapter describes schemes for adaptive routing and for adaptive wavelength assignment in an all-optical network, and examines the improvements these two kinds of adaptivity can offer on performance. For the specific adaptive routing scheme that we examine, we demonstrate that, for the hypercube and torus topologies considered, providing (at most) one alternate link at every hop gives a per-wavelength throughput that is close to that achieved by oblivious routing with twice the number of wavelengths per link. Also, we examine the effect of limited wavelength conversion in network performance and find that limited conversion to only one or two adjacent wavelengths can provide a considerable fraction of the improvement that full-wavelength conversion provides over no-wavelength conversion. These results clearly emphasize the need for network designers to investigate the tradeoffs between wavelength conversion, routing flexibility, and hardware cost when designing future optical networks.

The recent advances in fiber-optic technology are strongly affecting our everyday life and habits leading the way to the so-called "information technology society." The ever-increasing demand for bandwidth dictates imperatively the use of techniques and protocols that can optimally exploit the fiber's potential capabilities. One of the most promising technologies in this direction is *wavelength division multiplexing* (WDM), which is the current favorite multiplexing technology for long-haul communications in optical communication networks.[1] WDM divides the huge fiber's bandwidth (about 50Tbps) into many nonoverlapping WDM channels, each corresponding to a different wavelength (so as not to interfere with one another). With each WDM channel assigned to a different communication channel operating at (potentially) peak electronic rate (e.g., 10Gbps), WDM manages to accommodate multiple communication channels from different users (with dissimilar data formats) in parallel using the same fiber.

Emerging wide-area-networks that employ WDM are capable of switching data entirely in the optical domain by means of optical wavelength routing switches (WRSs) or wavelength selective cross-connects (WSXCs). In this way, all-optical connections, or *lightpaths*, which may span multiple fiber links, can be established across a network without undergoing any intermediate optical-to-electronic-to-optical (O-E-O) conversion. Figure 12.1 illustrates a wavelength-routed optical WDM network consisting of a set of nodes and a set of links. The network nodes can be either optical switching

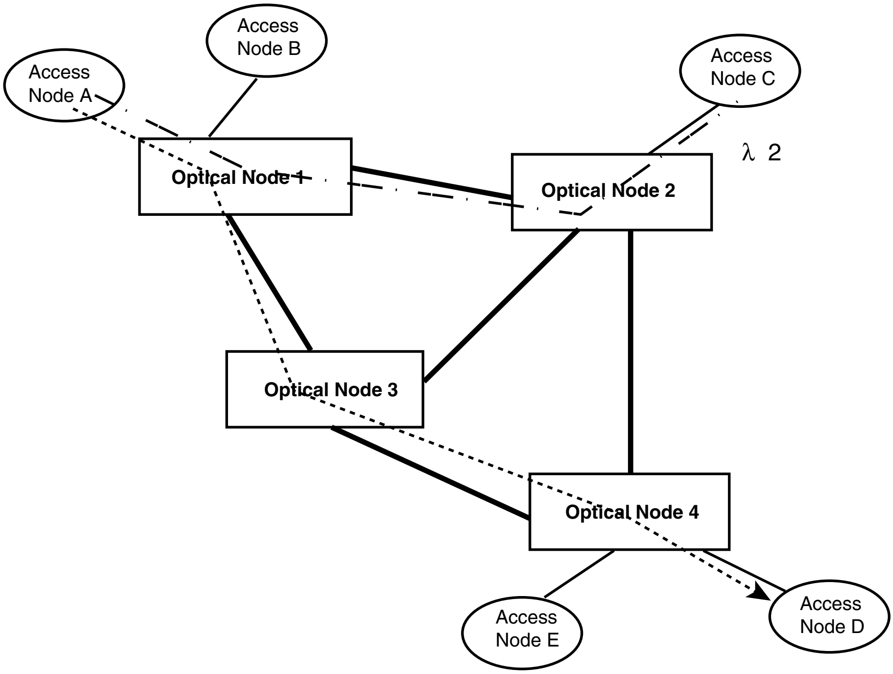

Figure 12.1 A wavelength-routed optical WDM network including optical switching nodes and access nodes. Lightpaths are established between pairs of access nodes on different wavelengths.

nodes or access nodes, while the links that join the nodes can be either inter-nodal links (connecting optical switching nodes, set in bold in Figure 12.1) or access links (connecting the access nodes to the optical network switches). The access nodes provide the necessary electronic-to-optical (and vice versa) conversion to interface the optical network to the conventional networks; they are connected to the optical network switches at a specific input fiber and wavelength. If a different wavelength is required, then wavelength conversion has to be used; the same is true for the traffic absorbed by an access station. Lightpaths can be established in this network between pairs of access nodes on different wavelengths.

A *unidirectional lightlink* between two switches is a specific wavelength within a specific link. A lightpath in an optical network is considered as a sequence of unidirectional *lightlinks* from a starting node *s* to a terminating node *d*. A lightpath consists of a specific wavelength of the input fiber in switch *s*, a specific wavelength of the output fiber in *d*, and the intermediate lightlinks that are used to connect the intermediate switches. The *physical length* of a lightlink is defined as the length of the corresponding fiber. The *amplifier length* of a lightlink is the number of amplifiers on the lightlink. The physical and the amplifier length of a lightpath are defined as the sums of the corresponding physical and the amplifier lengths of the lightlinks that comprise it. Dispersion and other factors pose an upper limit on the maximum allowable physical length of a lightpath. An upper limit also exists on the amplifier length of a lightpath (before undergoing O-E-O conversion). Note that the amplifier length of a lightpath is usually proportional to its physical length (because amplifiers are usually placed at equal distance), in which case the two types of constraints could be treated as one.

An optical switching device can be equipped, apart from amplifiers, with wavelength-conversion facilities, which enable switching data from a wavelength λ_i on an incoming link to a wavelength λ_j (j ? i) on an outgoing link. This conversion may be required for a connection arriving from some optical node that is switched to some other optical node (continuing connection), for a connection arriving from an access node (originating connection), or for a connection switched to an access node (terminating connection). In this context, we distinguish three classes of wavelength routing switches:

Author: Is this correct?

1. Switches with full-wavelength conversion capability, which can switch an incoming wavelength to *any* outgoing wavelength
2. Switches with limited-wavelength conversion capability, which can switch an incoming wavelength to a *subset* of the outgoing wavelengths
3. Switches with no-wavelength conversion capability, which can switch each incoming wavelength *only* to the same outgoing wavelength.

In the last case, a lightpath is required to occupy the same wavelength on each fiber link along its path in the network, a restriction known as the *wavelength-continuity constraint*, which increases the probability of call blocking. At each node, there is an upper limit on the number of O-E-O conversions that can be performed at that node. If wavelength conversion is also performed using O-E-O conversion (as opposed to using all-optical wavelength converters), then there is a single constraint for the total number of O-E-O conversions and wavelength conversions that can be performed at that node.

Full-wavelength conversion has been extensively studied and, although it has been shown to dramatically decrease the blocking probability and improve network performance,[7,8,11] it is not often feasible to provide such a capability to all the optical switches of a network, due to technological and financial limitations; wavelength conversion technology is still immature and expensive. These restrictions led researchers to some more practical alternatives: *sparse* and *limited* wavelength conversion. In sparse conversion, only *a few* of the switches have (full) conversion capabilities and, in this case, the objective is to minimize the number of such switches. Subramaniam et al. demonstrated that no significant degradation occurs in network performance when sparse — instead of full — wavelength conversion is employed.[19] In limited conversion, on the other hand, each switch provides wavelength conversion but with limited capabilities.

Author: Ref. 1 appears on page 2, but Refs. 2 through 6 are skipped.

Author: Should be a minus sign?

For example, in *k-adjacent wavelength switching*, an incoming wavelength can be translated only to a subset (called the *feasible wavelength set*) consisting of k of the W outgoing wavelengths (i.e., to k - 1 wavelengths in addition to itself). Limited conversion has received particular attention: Yates et al. presented a simple, approximate probabilistic analysis for single paths in isolation;[20] Ramaswami and Sasaki provided a non-probabilistic analysis for ring networks and, under certain restrictions, for tree networks and networks of arbitrary topology;[21] while Wauters and Demeester provided new upper bounds on the wavelength requirements for a WDM network under a static model of the network load.[22]

In an all-optical WDM network, connection establishment for a session involves two phases:

1. The selection of a route or sequence of hops the session must traverse
2. For each hop along the route and according to the wavelength conversion capability of the corresponding switching node, the selection of a wavelength on which the session will be carried for that hop

The path of a session is the sequence of link-wavelength pairs traversed by it. Path selection, therefore, involves *routing* and *wavelength assignment*, both of which may be either oblivious or adaptive. In oblivious (or static) routing, the route is selected at the source and is independent of the state (loading or congestion) of the network, while in adaptive (or dynamic) routing, the route is selected either at the source or on a hop-by-hop basis, based on the state of the network at the time of connection establishment.

Two well-known examples of static routing algorithms are *fixed routing* and *fixed-alternate routing*. In fixed routing, a single fixed path is predetermined for each source-destination pair, and a connection request is blocked if the associated path is not available. In fixed-alternate routing, a fixed set of predetermined paths is assigned to each source-destination pair. When a connection request arrives at a node, the set of predetermined paths is searched according to some policy (e.g., in fixed or adaptive order) to select an available path. Though static algorithms are much simpler to implement than adaptive routing schemes, they can lead to high blocking probabilities for connection requests.

In adaptive routing schemes, on the other hand, the path for a source-destination pair is selected dynamically by taking into account the network state, thus resulting in lower connection blocking probabilities. One form of adaptive routing is adaptive least-cost-path routing based on global network state information, in which every link is assigned a cost (based, for example, on wavelength availability) and, upon arrival of a connection request, a least-cost routing algorithm determines the path for the given source-destination pair. This scheme may be implemented in either a centralized or a distributed manner.[2] In the former version, a centralized entity maintains global network state information and establishes lightpaths in response to connection requests. In the distributed version, two common approaches are used: the *link-state* approach, where each node in the network maintains global network state information and, therefore, can find suitable routes for connection requests in a distributed manner;[3] and the *distance-vector* approach, where each node maintains a routing table that only indicates, for each destination and on each wavelength, the next hop to the destination and the distance to the destination.[4] A disadvantage of all the preceding adaptive routing schemes is the need for continuous network updates of the routing tables at each node, whenever network changes take place, which results in a significant increase in control overhead and a requirement for elaborate control and management protocols. Therefore, global knowledge-based adaptive routing schemes are mostly used in networks where lightpaths are quite static and do not change much with time.[2]

An important type of adaptive routing is *deflection routing*. Deflection routing protocols have previously been analyzed by several researchers under a variety of assumptions on the underlying network topology, and have been shown to perform outstandingly in many cases due to their low overhead and high adaptivity.[23, 25–33] The deflection routing schemes proposed to date, however, are based on packet-by-packet (datagram) deflections, and may be inappropriate for high-speed networks due to their excessive per-packet processing requirements, the loss of packet order, etc. Thus, this chapter concentrates on a form of deflection routing called the *virtual circuit deflection* (VCD) protocol, first proposed in Varvarigos and Lang, which performs deflections on a per-session (virtual circuit) basis.[17] VCD protocol is a hybrid of virtual circuit switching and deflection routing, combining some of their individual advantages. It alleviates to a large extent many of the problems of previous datagram deflections schemes, while its small buffer requirements make it particularly appropriate for high-speed networks that use optical switching.

The VCD scheme examined in this chapter is a virtual circuit switching protocol of the tell-and-go variety, where data starts being transmitted shortly after the setup packet of a session is sent. In this scheme, the intermediate links (and wavelengths) of a path are determined dynamically on a hop-by-hop (instead of end-to-end) basis, depending on link (and wavelength) utilization. At each node, an outgoing link is selected from among the subset of outgoing links that lie on a shortest route to the destination. If wavelength resources are unavailable on the chosen link, an alternate link lying on the shortest route to the destination is tried; we then say that the session is deflected. This process continues until either an available link is found or all the alternate links have been examined. Hence, routing-table updates in the network are not needed, and control overhead is greatly reduced.

Wavelength assignment usually does not take place in parallel with the selection of the links of the path, as we assume it happens with the VCD scheme; instead, once a route has been selected for a source-destination pair, a wavelength assignment algorithm assigns suitable wavelengths to each link of the route, so that any two lightpaths sharing the same physical link are assigned different wavelengths. In the static case (i.e., when the lightpaths that are to be set up are known in advance), and under the wavelength continuity constraint discussed before, wavelength assignment reduces to the *graph-coloring problem,* which is known to be NP-complete. Heuristic methods (such as random assignment, first-fit, least-used assignment, etc.) are usually employed to assign wavelengths to lightpaths. For a review of these methods and for performance comparison in terms of connection blocking, see Zang et al.[5]

Amy: Spell out NP?

The design of efficient routing and wavelength assignment (RWA) algorithms in all-optical networks has been the objective of many research initiatives. Karasan and Ayanoglu analyzed the first-fit wavelength assignment strategy in a network with no-wavelength conversion and fixed shortest-hop routing.[14] They also proposed an adaptive RWA algorithm and evaluated its performance via simulations. Mokhtar and Azizoglu also proposed several adaptive RWA algorithms for networks with no-wavelength conversion; they also analyzed oblivious alternate routing using a fixed-order wavelength search.[15] Harai et al. analyzed oblivious alternate routing with fixed wavelength assignment and no-wavelength conversion.[16] Harai et al. also analyzed oblivious alternate routing with various wavelength assignment schemes for networks with limited wavelength conversion.[10] All these algorithms, however, require information on global wavelength utilization, assuming either a periodic exchange of such information[15] or a centralized network controller.[14,16]

The remainder of this chapter is organized as follows. Section 12.2 describes the virtual circuit deflection (VCD) scheme and shows how it can be combined with other techniques to improve network performance. Section 12.3 presents analytical results for oblivious and adaptive routing and wavelength assignment in the torus and hypercube networks with full-wavelength conversion. Section 12.4 focuses on the simulation results obtained for the VCD scheme; furthermore, we examine the effects of wavelength conversion in network performance as well as present performance results for the VCD protocol, focusing mainly on the adaptivity VCD can exhibit. Also, some interesting design options when building an all-optical network are discussed, and Section 12.5 concludes the chapter.

12.2 *Virtual circuit deflection (VCD): an adaptive routing scheme*

This section describes a specific adaptive routing scheme — VCD. We assume that connection requests are generated at the source nodes with a specified destination and bandwidth requirement (number of wavelengths required). A source node tries to accommodate each request by choosing one of its outgoing links that lies on a shortest route to the destination (according to a static routing table, which is held at each node* and is based on the [possibly outdated] network topology). If the chosen link does not have the

Amy: Footnote needs to be moved from page 8 to page 7.

wavelength resources required for the connection, an alternate link laying on the shortest route to the destination is tried (deflection) until either an available link is found or all the alternate links have been examined. After determining an outgoing link, a setup packet is transmitted to the next node of the path to set the routing tables and reserve resources at intermediate nodes. At each hop, the setup packet randomly selects a wavelength from among the available wavelengths of a link and, if it is successful in establishing a connection, the wavelengths required by the session are reserved for the session duration; otherwise, the session is randomly assigned a new time at which to try. The setup packet is thus forwarded hop-to-hop and is followed after a short delay by the data packets. If the setup packet is successful in reserving resources on all the links on the path to the destination without deflection, the VCD scheme looks like the usual (forward) reservation protocols, with the difference that the reservation (setup) phase and the transmission phase overlap in time.

In a large mesh, most intermediate nodes have two outgoing links lying on a shortest route. In a hypercube, i outgoing links are lying on a shortest route when the packet is at a distance i from its destination. The number of alternate links that lie on a shortest route to the destination may change as the setup packet progresses toward its destination. Furthermore, a limit on the number of alternate links that are examined could be used to reduce further congestion or the processing overhead at an intermediate node. We let l be the number of outgoing links that a session may try at each hop, which we refer to as the routing flexibility. The number of *feasible* outgoing links at a node t is given by $\min(l, n_{t,d})$, where $n_{t,d}$ is the number of outgoing links at node t that lie on a shortest route to the destination d. Therefore, if the capacity of each link is divided in k wavelengths, a session currently at node t will be blocked and scheduled to retry only if all of the k wavelengths on each of its feasible outgoing links are unavailable.

It is possible for sessions to be deflected such that the paths contain loops. This may arise after a series of deflections or if a setup packet is deflected immediately to the previously visited node. In either case, the resources reserved in the loop are inefficiently used and it is desirable to remove the loop; however, unless the setup packet visits the intermediate node for the second time prior to the arrival of the first data packet, it is unclear whether the added protocol complexity associated with removing the loop outweighs the efficiency benefits.

Allowing sessions to follow very long paths can waste network resources, increasing the probability that future sessions will be blocked or forced to take even longer paths. To avoid the waste that occurs when a session follows a very long path due to deflections, we may request that a session is dropped when the setup packet has traveled more than H hops without reaching its destination. The parameter H can be chosen to be equal to a multiple (e.g., two or three times) of the shortest distance between the source and the destination of the session, and it may also be dependent on the current congestion in the network. A session that has undergone too many deflections is dropped by transmitting a control packet to the source, requesting it to cease transmitting new packets. . Data packets sent prior to the arrival of the control packet at the source can either be dropped or allowed to remain in the network until they reach their destination (possibly over a very long path), while the remaining data is sent later, over a different path, by the source.

The time gap between the transmission of the setup packet and the transmission of the first data packet from a source is chosen to be equal to the maximum number of hops *H* allowed for the particular session times the processing time of a setup packet at a node. In other words, the gap must be at least as large as the minimum time by which the connection setup phase and the data transmission phase should be separated in order to ensure that data packets do not overpass the setup packet.

Author: Please clarify.

Author: Amy: Correct word?

12.2.1 *Comparison with wait-for-reservation and tell-and-go protocols*

A sizable portion of the traffic in future multi-gigabit networks will involve the high-speed transfer of massive amounts of data at nearly constant rates, and will require guaranteed lossless delivery and explicit resource reservation (e.g., the constant-bit rate class of asynchronous transfer method (ATM) traffic). Most of the connection control protocols designed to deal with this type of traffic use explicit reservations prior to the transmission of any data. Because a source has to wait for an acknowledgment from the destination in such protocols before it can transmit any data packets, we refer to them as wait-for-reservation virtual circuit (WRVC) protocols. WRVC protocols tend to be inefficient in terms of link utilization because network resources are reserved for more time than a session requires. Furthermore, the pre-transmission delay required for the setup phase is often significant compared with the delay requirements of the session and unwarranted if the network load is light.

For sessions in which the round-trip pre-transmission delay is not acceptable, *tell-and-go* protocols are more appropriate. In such protocols, the setup packet is followed after a short delay by the data packets, achieving in this way a pipelining between the setup phase and the data transmission phase as well as reducing the pre-transmission delay to the minimum possible. If the available resources found by the setup packet at an intermediate node are not adequate to accommodate the session, the excess data packets are usually buffered at the node, and backpressure[35] is exercised to upstream nodes to control the source transmission rate. The VCD scheme avoids the difficulties associated with wait-for-reservation and backpressure-based protocols, and can ensure lossless communication with little buffering and a small pre-transmission delay.[17] Because link wavelength resources are reserved for a duration that is slightly larger than the holding time of a session and are available for the remaining time, VCD has an efficiency advantage over WRVC protocols. This is particularly important for high-speed networks where propagation times are often large compared with the typical holding time of sessions.

12.3 Performance analysis: flexibility in routing vs. flexibility in wavelength assignment

This section presents analytical results for oblivious and adaptive routing and wavelength assignment in the *torus* and *hypercube* networks with full-wavelength conversion. Our choice of the torus and hypercube topologies reflects our interest in analyzing two popular topologies with very different characteristics. The torus is a sparse topology with a small (fixed) node degree and rather large diameter, while the hypercube is a dense topology, with a node degree and diameter that increase logarithmically with the number of nodes. The results are based on the analysis found in Lang et al., where the reader is referred for a more in-depth study.[6] The analytical results apply to regular, all-optical networks with full-wavelength conversion. These results hold for any vertex and edge-symmetric topology and, with modifications, to any vertex symmetric (but not edge-symmetric) topology.

We assume a distributed network model where the routing decision is made locally at each node, using information only about the state of each node's outgoing links and wavelengths. Also, we do not require that the alternate paths between a source-destination pair be link disjoint,[15,16] instead allowing links (and wavelengths) to overlap between alternate paths. In the network model considered, new sessions with uniformly distributed over all nodes destinations arrive independently at each node of the network according to a Poisson process. The capacity of each link is divided into k wavelengths, and each node has full-wavelength conversion capability. An outgoing link of a node with k wavelengths per link is modeled by an auxiliary $M/M/k/k$ queuing system. Using the occupancy distribution of this system, a closed-form expression for the probability P_{succ} of successfully establishing a circuit can be produced without the need to use the link independence blocking assumption, but instead by taking into account partially the dependence between the acquisition of successive wavelengths on the path followed by a session. The analysis presented is general, computationally inexpensive, and scales easily for larger network sizes and arbitrary k. It applies to both oblivious and adaptive routing, and applies equally well to multi-fiber networks with no-wavelength conversion.[6]

We examine how the extent of improvement in achievable throughput for a fixed P_{succ} depends on the *number of wavelengths k per link* and on the number of links l that may be tried at each hop. This is important because it impacts on the cost and the complexity of the switch. Increasing the *routing flexibility l* increases the switch complexity and delay. Similarly, with full-wavelength conversion, increasing the number of wavelengths k per link increases hardware complexity and may be difficult to realize with current technology. We find that although the throughput per wavelength increases superlinearly with k, the incremental gain in throughput per wavelength (for a fixed P_{succ}) saturates rather quickly to a linear increase. We also see that when the routing flexibility l is varied, the largest incremental gain in throughput per wavelength occurs when l is increased from one to two. We also compare the performance obtainable with a certain number of wavelengths k with that obtainable with a certain routing flexibility l. For the torus and hypercube topologies, we find that for a fixed P_{succ}, a system with k wavelengths per link and only one alternate choice of an outgoing link (i.e., $l = 2$) gives a per-wavelength throughput that is close to that achieved by a system using oblivious routing with $2k$ wavelengths per link, with only a small additional improvement as l is increased further. The preceding observations imply several interesting alternatives for the provisioning and expansion of all-optical networks, some of which we discuss in Section 12.4.

12.3.1 Torus networks

We consider the $p \times p$ torus network, which consists of $N = p^2$ nodes arranged along the points of a two-dimensional (2-D) grid with integer coordinates, with p nodes along each dimension. Two nodes (x_1, x_2) and (y_1, y_2) are connected by a bidirectional link if and only if, for some $i = 1, 2$, we have $(x_1 - y_i) \bmod p = 1$ and $x_j = y_j$ for j ? i. In addition to these links, wraparound links connecting node $(x_1, 1)$ with node (x_1, p), and node $(1, x_2)$ with node (p, x_2), are also present.

Author: Is this correct?

In oblivious routing with full-wavelength conversion, the route followed by a session is chosen at the source and is independent of the state of the links. In this case, a session is blocked and scheduled to retry only if all wavelengths on the desired outgoing link are unavailable, where we assume that a setup packet selects the outgoing wavelength from among the available wavelengths on the link with equal probability. We consider an X - Y routing scheme where a session follows a shortest route to its destination, first traversing all the links in one dimension (horizontal or vertical) and then traversing all the links in the other dimension (vertical or horizontal); the first dimension is selected as random at the source. For uniformly distributed destinations, the average probability of success for a new session can be calculated to be:

Author: Should be a minus sign?

$$P_{succ} = \begin{cases} \dfrac{\delta}{\alpha_1}\left[\left(1 + 2\alpha_1\left(\dfrac{1-\alpha_2^{(p-1)/2}}{1-a_2}\right)\right)^2 - 1\right] & p \text{ odd} \\[2ex] \dfrac{\delta}{\alpha_1}\left[\left(1 + 2\alpha_1\left(\dfrac{1-\alpha_2^{p/2}}{1-a_2}\right) + a_1\left(\dfrac{1-\alpha_2^{(p/2)-1}}{1-a_2}\right)\right)^2 - 1\right] & p \text{ even} \end{cases} \qquad (12.1)$$

where $\delta = \alpha_0/(p^2 - 1)$, α_0 expresses the probability that a wavelength is available on the outgoing link of the originating node of the session, and α_1 and α_2 express the probability that a wavelength is available on an outgoing link at a transit node of the path, given that a wavelength was available on an incoming link of that node.

In adaptive VCD routing, a link is selected at random, at each hop, from among all the outgoing links that lie on a shortest route to the destination; if all the wavelengths on the chosen link are unavailable, an alternate link lying on the shortest route to the destination is tried. This process continues until either an available link is found or all the alternate links have been examined. For the torus network, there are at most two outgoing links at a node that lie on a shortest route to the destination. For uniformly distributed destinations, the average probability of success for a new session P_{succ} can be calculated to be:

$$P_{succ} = \begin{cases} \dfrac{1}{p^2-1}\left[\displaystyle\sum_{i-1}^{(p-1)/2} 4iP_{succ}(i) + \sum_{i=(p+1)/2}^{p-1} 4(p-i)P_{succ}(i)\right] & p \text{ odd} \\ \dfrac{1}{p^2-1}\left[\begin{array}{l}\displaystyle\sum_{i-1}^{(p/2)-1} 4iP_{succ}(i) \\ \displaystyle + \sum_{i=(p/2)+1}^{p-1} 4(p-i)P_{succ}(i) \\ +2(p-1)P_{succ}\left(\dfrac{p}{2}\right) + P_{succ}(p)\end{array}\right] & p \text{ even} \end{cases} \quad (12.2)$$

where $P_{succ(i)}$ is defined as:

$$P_{succ}(i) = \begin{cases} (a_0(1)+(i-1)a_0(2))\left(\dfrac{1}{i!}\right) \\ \cdot\displaystyle\prod_{b=1}^{i-1}(a_1(1)+(b-1)a_1(2)), & i \le \left[\dfrac{p-1}{2}\right] \\ a_0(2)a_1(2)^{i-[(p+1)/2]}\left(\dfrac{1}{([(p-1)/2])!}\right) \\ \cdot\displaystyle\prod_{b=1}^{[(p-1)/2]}(a_1(1)+(b-1)a_1(2)), & i > \left[\dfrac{p-1}{2}\right] \end{cases} \quad (12.3)$$

where $\alpha_0(i)$, $i = 1, 2$ expresses the probability that at least one wavelength on $i = \min(l, n_{s,d})$ outgoing links at the origin is available, and $\alpha_1(i)$, $i = 1, 2$ expresses the probability that, at each transit node t, a wavelength is available on one of $i = \min(l, n_{t,d})$ alternate outgoing links, given that a wavelength was available on an incoming link of that node.

12.3.2 Hypercube networks

This section considers the 2^r-node hypercube network, where each node can be represented by a binary string $(x_1, x_2, \ldots, x_r)$, and two nodes are connected via a bidirectional link if their binary representations differ in only one bit.

In oblivious routing, where a shortest route is chosen at random at the source, and for uniformly distributed destinations, the average probability of success for a new session can be calculated to be:

$$P_{succ} = \frac{a_0}{a_1\left(2^r - 1\right)}\left[\left(1 + a_1\right)^r - 1\right] \tag{12.4}$$

where α_0 and α_1 are as defined in Eq. (12.1).

In adaptive VCD routing we note that in the hypercube network, a node that is i hops away from the destination has i outgoing links lying along a shortest route to the destination. We let l, $l = r$, be the maximum number of outgoing links that may be tried at any hop. Assuming the source is at a distance i hops from its destination, the probability of successfully establishing a connection is given by:

Author: Is this correct?

$$P_{succ}(i) = a_0\left(\min(l, i)\right)\prod_{j=1}^{i-1} a_1\left(\min(l, j)\right) \tag{12.5}$$

where α_0 $(\min(l, i))$, $i = 1, \ldots, r$ and $\alpha_1(\min(l, j))$, $j = 1,.., r$ are as defined in Eq. (12.3). Then, the probability P_{succ} can be found using Eq. (12.4).

12.4 Simulation results

Before we examine the simulation results for the VCD scheme and compare them with the analytical results presented in Section 12.3, it is worth turning our attention first to the effects of a number of parameters on the performance of the network under study.

First, we examine the effect of wavelength conversion in network performance, considering the torus network with oblivious routing and in three different cases:

1. No-wavelength conversion (or 1-adjacent wavelength switching)
2. Limited-wavelength conversion using k–adjacent wavelength switching, where $k = 2, 3$

3. Full-wavelength conversion (or W-adjacent wavelength switching) in a WDM network with W wavelengths per link.

We note that full-wavelength conversion provides the best achievable performance (in terms of the realizable probability of success for a given arrival rate per wavelength or in terms of the realizable throughput per wavelength for a given probability of success) for a given number of wavelengths W per link. When no-wavelength conversion is used, the different wavelengths on a link do not interact with one another. Thus, an all-optical network with W wavelengths per fiber is essentially equivalent to W disjoint single-wavelength networks operating in parallel. To obtain the probability of success in this special case, it is therefore enough to focus attention on any one of the W independent parallel networks, for which the analysis given in Sharma applies.[34]

We define the degree of conversion δ of a k-adjacent wavelength switching system with W wavelengths per fiber to be

$$\delta = \frac{k-1}{W-1} \times 100\%$$

Thus, $\delta = 100\%$ corresponds to the case of full-wavelength conversion (or W-adjacent wavelength switching), while $\delta = 0\%$ corresponds to the case of no-wavelength conversion (or 1-adjacent wavelength switching).

We define $P_{succ}(\lambda, k)$ to be the probability of success in a k-adjacent wavelength switching system when the arrival rate λ per node per wavelength is equal to v/W; and we define $\lambda(P_{succ}, k)$ to be the throughput per node per wavelength of a k-adjacent wavelength switching system, when the probability of success is equal to P_{succ}. To quantify the performance of limited wavelength conversion vs. full- or no-wavelength conversion, we also define the *throughput efficiency* $\lambda(P_{succ}, k)$ of a k-adjacent wavelength-switching scheme, with W wavelengths per fiber, for a given probability of success P_{succ}, to be:

$$\Delta\lambda(P_{succ}, k) = \frac{\lambda(P_{succ}, k) - \lambda(P_{succ}, 1)}{\lambda(P_{succ}, W) - \lambda(P_{succ}, 1)} \times 100\%$$

and the *success efficiency*$P_{succ}(\lambda, k)$ of a k-adjacent wavelength switching system, for a given arrival rate per node per wavelength λ, to be:

$$\Delta P_{succ}(\lambda, k) = \frac{P_{succ}(\lambda, k) - P_{succ}(\lambda, 1)}{P_{succ}(\lambda, W) - P_{succ}(\lambda, 1)} \times 100\%$$

The throughput and success efficiencies represent the degree of improvement (over no-wavelength conversion) in the throughput and in

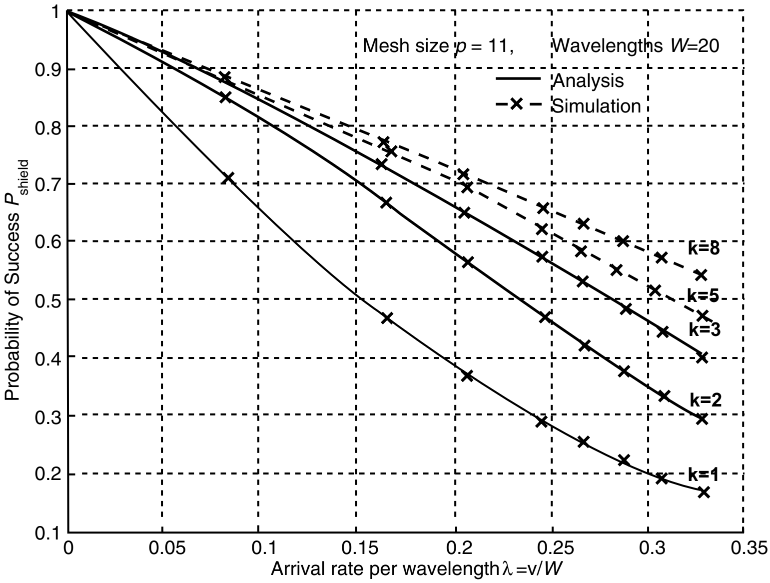

Figure 12.2 Success probability P_{succ} vs. the arrival rate per wavelength λ, for a $p \times p$ torus (p = 11), for W = 8 wavelengths per link.

the probability of success respectively, which is obtained when limited wavelength conversion with k-adjacent wavelength switching is used, as a percentage of the improvement obtained when full-wavelength conversion is used. For $k = W$ (full-wavelength conversion), we get $\lambda(P_{succ}, k)$ = 100% and $P_{succ}(\lambda, k)$ = 100%, while for $k = 1$ (no-wavelength conversion), we get $\lambda(P_{succ}, k)$ = 0% and $P_{succ}(\lambda, k)$ = 0% (no improvement).

In Figure 12.2, we present performance results for the probability of success P_{succ} plotted vs. the arrival rate per node per wavelength $\lambda = v/W$ when limited wavelength conversion to only one or two additional wavelengths (i.e., k = 2, 3) is permitted. The results depicted here were obtained using the analysis presented in Sharma and Varvarigos, and Sharma.[9,34] Observe that limited conversion to only one or two adjacent wavelengths provides a considerable fraction of the improvement that full-wavelength conversion provides over no-wavelength conversion. These benefits are summarized in Table 12.1, where we illustrate the throughput and success efficiencies for a $p \times p$ torus (p = 11) for a few selected points.

Table 12.1Table 12.1 Quantifying the Benefits Obtained with Limited Wavelength Conversion when K = 2, 3

p	W	?P_{succ}(0.25, 2)	?P_{succ}(0.25, 3)	?λ(0.7, 2)	?λ (0.7, 3)
11	5	51%	86%	61%	87%
11	20	47%	79%	49%	76%

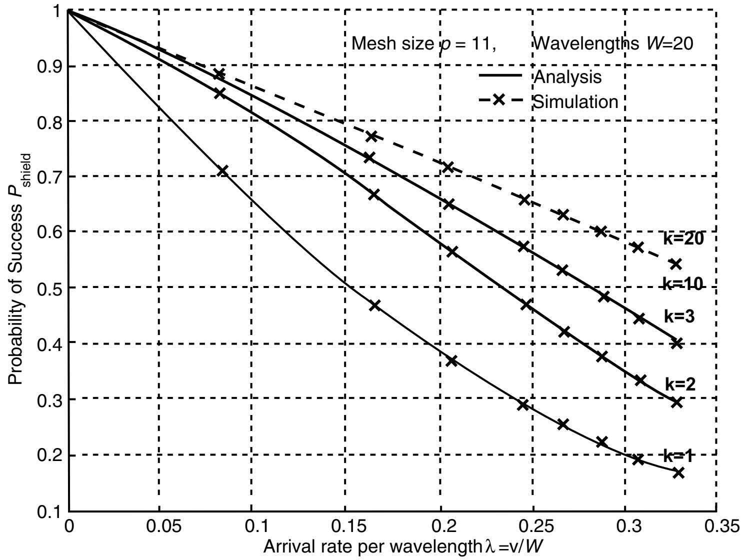

Figure 12.3 Success probability P_{succ} vs. the arrival rate per wavelength λ, for a $p \times p$ torus (p = 11), for W = 20 wavelengths per link.

Also, the benefits of wavelength conversion diminish as the extent of conversion k increases and, eventually, appear to saturate. We see, therefore, that limited conversion of small range (i.e., k = 2 or 3) gives most of the benefits obtained by full-wavelength conversion, where $k = W$. For instance, in Figure 12.3, which also illustrates the network performance for $k = W = 20$ wavelengths, increasing the extent of conversion k beyond some value leads to diminishing returns. Similar remarks regarding the effects of the extent of wavelength conversion also apply in the case of hypercube networks, using either descending dimensions switches or crossbar switches.[9,34]

Next, we present performance results for the VCD protocol, focusing mainly on the adaptivity that VCD can exhibit. The results are obtained from Varvarigos and Lang, where a Manhattan street (MS) network topology is considered.[17] The MS network is a two-connected regular mesh network with unidirectional communication links, which has been analyzed extensively in the literature for datagram deflection schemes due to its regularity and symmetry properties.[24,26,30,31] The MS $X \times Y$ -dimensional wraparound mesh consists of $N = XY$ processors arranged along the points of a 2-D space that have integer coordinates. X processors exist along the x-dimension, and Y processors exist along the y-dimension, where X and Y are even numbers. Each processor has two outgoing links, one horizontal and one vertical. The horizontal links are directed eastward on even rows and westward on odd rows, while the vertical links are directed northward on even columns and southward on odd columns. Each processor is represented by a pair (x, y) with $0 \le x \le X - 1$ and $0 \le y \le Y - 1$.

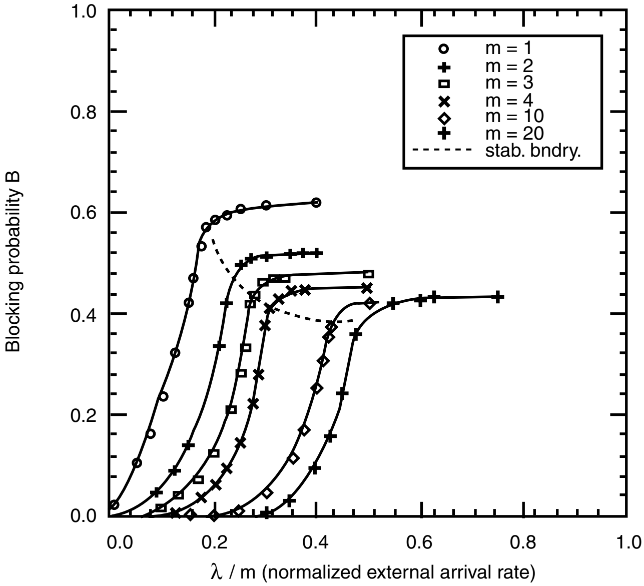

Figure 12.4 Performance results for the VCD protocol, illustrating the blocking probability *B* as a function of the normalized arrival rate per unit of link capacity λ/m for an 8×8 MS network and several values of *m*. The dashed lines correspond to the stability boundary.

A natural measure of the performance of the VCD protocol is the *inefficiency ratio* $\eta(\lambda)$, defined as the ratio

$$\eta(\lambda) = \frac{D(\lambda)}{D(0)}$$

of the average path length $D(\lambda)$ taken by a session for a given arrival rate , over the average shortest-path length $D(0)$ of the MS network topology. The inefficiency ratio characterizes the effectiveness with which the VCD protocol uses the network bandwidth for a given network load. In Figures 12.4 and 12.5, we illustrate the blocking probability *B* (the probability that a session attempting to establish a connection is blocked at its first hop) and the inefficiency ratio $\eta(\lambda)$, as a function of the normalized arrival rate per unit of link capacity λ/m, for an 8×8 MS network. In fact, *m* in this case can be viewed as the number of sessions or channels that can simultaneously use a link and hence, *m* provides a measure of the adaptivity of VCD. The dashed lines in these figures highlight the stability boundary (points to the left of the boundary belong to the stable region), where the stable region is defined as the region where the connection request queue remains finite; stability is not directly related to *B,* and it is possible to have *B* considerably less than one and still be in the unstable region. From Figures 12.4 and 12.5,

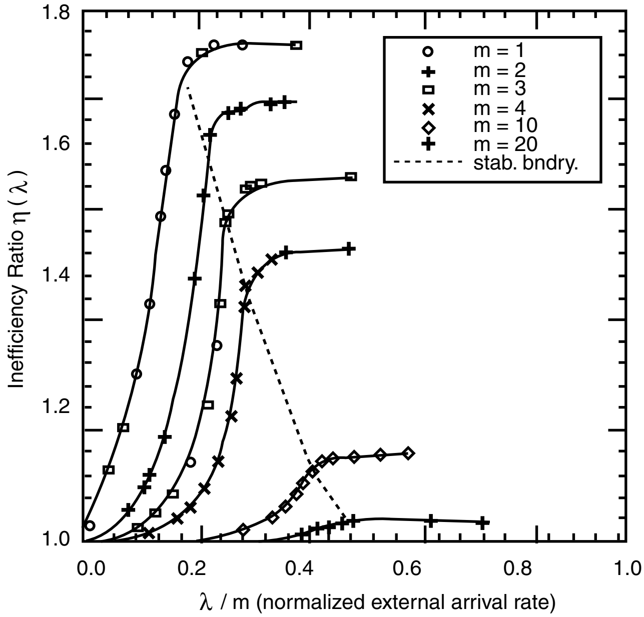

Figure 12.5 Performance results for the VCD protocol, illustrating the inefficiency ratio $\eta(\lambda)$ as a function of the normalized arrival rate per unit of link capacity λ/m, for an 8×8 MS network and several values of m. The dashed lines correspond to the stability boundary.

it is evident that when link capacity m is large (i.e., more channels can simultaneously use the link and, therefore VCD adaptivity is enhanced), the efficiency of the VCD protocol increases significantly. For example, for m = 20, the blocking probability B is always less than 0.4 and the lengths of the paths taken are, on the average, within 5% from the shortest path length for any value of the external arrival rate λ.

Now we will focus on the effect of full-wavelength conversion on oblivious and adaptive routing. In Figures 12.6 and 12.7, we present the success probability P_{succ} predicted by the analytical results in Section 12.3 vs. the results obtained from simulations for oblivious routing in the hypercube and torus networks, respectively; Figures 12.8 and 12.9 present the respective results for adaptive VCD routing. We observe that in all the figures, close agreement exists between the simulations and the analytically predicted values over the entire range of applicable input rates.[21] Despite its accuracy, the presented analysis is considerably simpler than the analyses available in the literature and its computational requirements are modest, allowing it to scale easily for large k.

To compare the performance of systems with varying k and l, we define the *incremental per-wavelength throughput gain* $\lambda(k_1, l_1; k_2, l_2)$ of a system with k_2 wavelengths and a choice of l_2 links per hop, over a system with k_1 wavelengths and a choice of l_1 links per hop, for a given P_{succ}, to be:

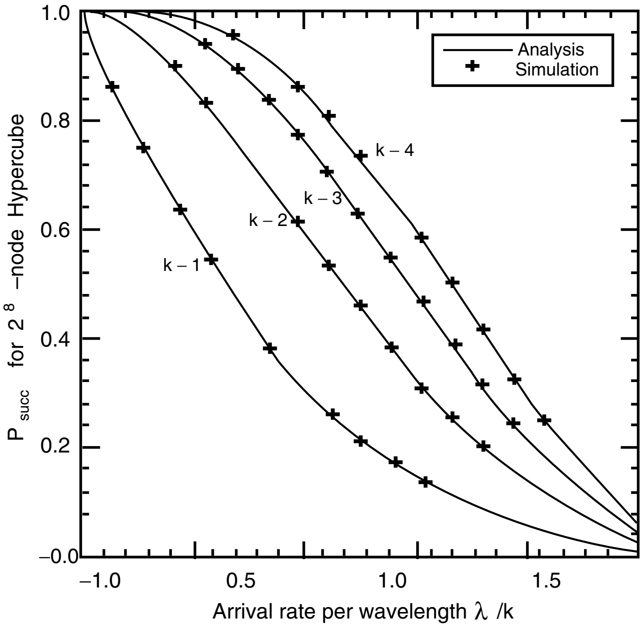

Figure 12.6 Analytical and simulation results for P_{succ} vs. the arrival rate per wavelength λ/k, for a2^6-node hypercube network, using oblivious routing.

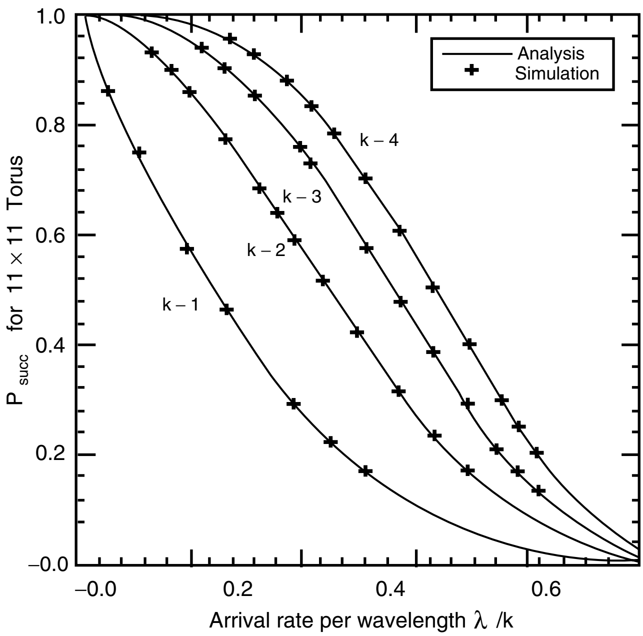

Figure 12.7 Analytical and simulation results for P_{succ} vs. the arrival rate per wavelength λ/k, for an11×11 torus network, using oblivious (X–Y) routing.

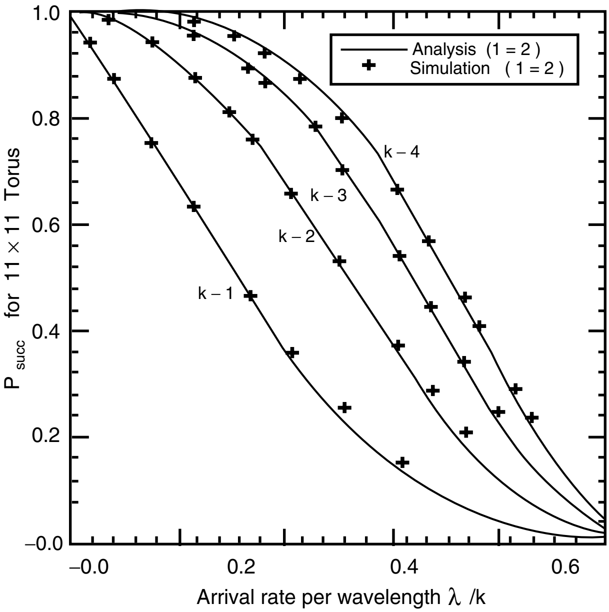

Figure 12.8 Analytical and simulation results for P_{succ} vs. the arrival rate per wavelength λ/k, for an 11 × 11 torus hypercube network, using adaptive VCD routing.

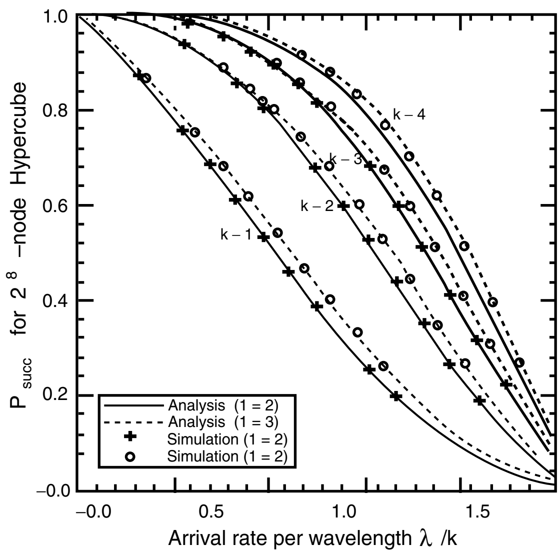

Figure 12.9 Analytical and simulation results for P_{succ} vs. the arrival rate per wavelength λ/k, for a2^6-node hypercube network, using adaptive VCD routing.

$$\Delta\lambda(k_1, l_1; k_2, l_2) = \frac{\lambda(P_{succ}, k_2, l_2) - \lambda(P_{succ}, k_1, l_1)}{\lambda(P_{succ}, k_1, l_1)} \times 100\% \qquad (12.6)$$

where $\lambda(P_{succ}, k, l)$ is the throughput per node per wavelength in a system with k wavelengths and routing flexibility l, when the probability of success is equal to P_{succ}.

We also define the *incremental probability of success gain* $P_{succ}(k_1, l_1; k_2, l_2)$ of a system with k_2 wavelengths and a choice of l_2 links per hop, over a system with k_1 wavelengths and a choice of l_1 links per hop, for a given λ/k, to be:

$$\Delta P_{succ}(k_1, l_1; k_2, l_2) = \frac{P_{succ}(\lambda, k_2, l_2) - P_{succ}(\lambda, k_1, l_1)}{P_{succ}(\lambda, k_1, l_1)} \times 100\% \qquad (12.7)$$

where $P_{succ}(\lambda, k, l)$ is the probability of success in a system with k wavelengths and routing flexibility l, when the probability of success is equal to λ/k.

The throughput and probability of success gains measure the degree of improvement that a full-wavelength conversion system with k_2 wavelengths and a choice of l_2 outgoing links per hop provides over a similar system with k_1 wavelengths and a choice of l_1 links per hop.

In Figures 12.10 and 12.11, we illustrate the analytically predicted probability of success P_{succ} vs. the arrival rate per wavelength λ/k, for k ranging from 1 to 16, for the torus and hypercube networks, respectively.

As can be seen from Figures 12.10 and 12.11, for a given P_{succ} and fixed l, the throughput per wavelength increases with increasing k. In other words, the throughput per link (and the network throughput) increases superlinearly with k. The linear part of the increase in throughput is because of the increase in capacity, while the superlinear part of the increase is due to more efficient use of that capacity because of the greater flexibility in establishing a circuit when a larger number of wavelengths is available. The incremental gain in achievable throughput per wavelength for a given l, $\lambda(k_1, l; k_2, l)$, however, decreases rapidly with increasing k. This result holds for both oblivious and adaptive VCD routing, and is in agreement with the results for oblivious routing presented in Sharma and Varvarigos, and Koch.[9,18] Similarly, the incremental throughput gain for a given k, $\lambda(k, l_1; k, l_2)$, decreases rapidly with increasing l. If we fix l_1 and l_2, and increase k, the incremental gain decreases, suggesting that the performance improvement for adaptive VCD routing is tightly coupled with the number of wavelengths, and that the benefits of alternate routing are not as significant when the number of wavelengths k is large.

Author: Underline correct?

Another interesting feature of adaptive VCD routing in networks with full-wavelength conversion is that the per-wavelength throughput for fixed number of wavelengths k and increasing the routing flexibility l, appears to saturate at or near the per-wavelength throughput of a system using oblivious routing with wavelength conversion over twice as many wavelengths.

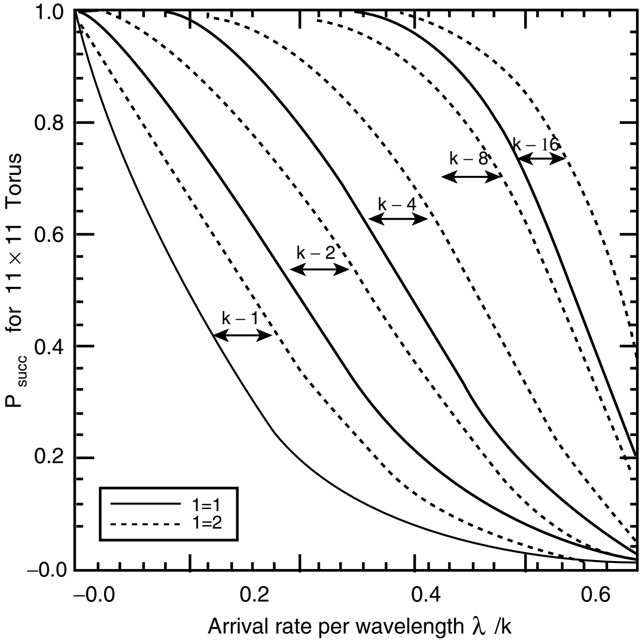

Figure 12.10 The probability of success P_{succ} for an 11 × 11 torus network, for *k* varying from 1 to 16.

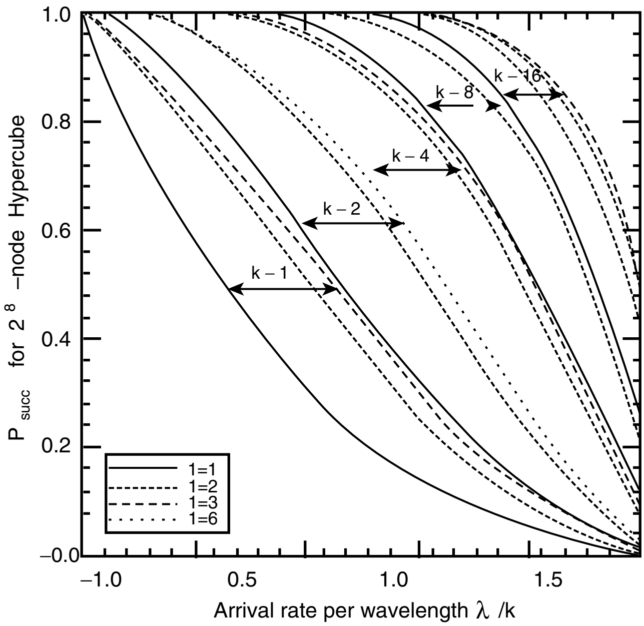

Figure 12.11 The probability of success P_{succ} for a 2^6-node hypercube network, for *k* varying from 1 to 16.

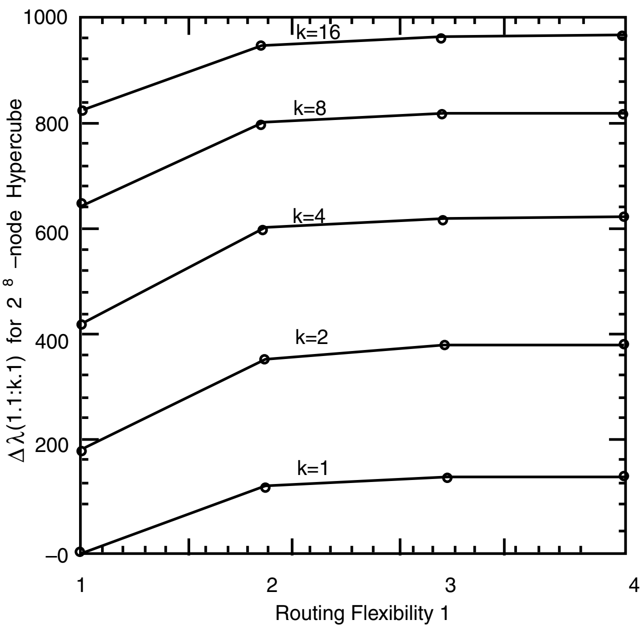

Figure 12.12 The incremental throughput gain λ*(1, 1; k, l)*, for *k* = 1, 2, 4, 8, 16 and *l* = 1, 2, 3, 4], for a 2^6-node hypercube network with P_{succ} = 0:8.

In Figure 12.12, we plot the incremental throughput gain for the hypercube network when P_{succ} = 0.8 and the number of wavelengths *k* ranges from 1 to 16, and the routing flexibility *l* ranges from 1 to 4 (i.e., we plot λ(*1, 1; k, l*) for *k* = 1, 2, 4, 8, 16 and *l* = 1, 2, 3, 4). As depicted in Figure 12.12, the largest increase in incremental throughput gain occurs when the routing flexibility increases from *l* = 1 to *l* = 2, regardless of the number of wavelengths. Furthermore, this gain obtained by increasing the routing flexibility from *l* = 1 to *l* = 2, with fixed *k*, approaches the gain obtained by doubling the number of wavelengths to 2*k*, with *l* = 1. For example, the incremental throughput gain for *k* = 8 and *l* = 2 is within 3% of the incremental throughput gain for *k* = 16 and *l* = 1.

The previous discussion leads to some interesting design options when building an all-optical network. For instance, because the per-wavelength throughput gain saturates quickly with increasing *k*, simply building a network in which every node can translate between *k* wavelengths may not be the most efficient option. Instead, it may be preferable to build a network in which every node consists of *k/n* simpler switching elements operating in parallel (each switching between a nonintersecting subset of *n* wavelengths) that achieves performance comparable to that of the *k*-wavelength system at a much lower cost. This suggests that a network designer may initially choose to build the network with nodes that have a small number of parallel channels, with *n* wavelengths per channel. As network traffic grows, the designer may expand the nodes by adding more parallel channels. Better yet, instead of increasing the number of channels per link at every network node, the designer may focus on the routing algorithms and may choose to increase the routing flexibility to obtain equivalent performance at no extra hardware cost. For instance, the designer may simply increase the number of outgoing links that may be tried at each hop. Observe, however, that the routing flexibility is limited by the network topology and is also a function of the switch architecture. Our results emphasize the need for network designers to investigate the tradeoffs between wavelength conversion, routing flexibility, and hardware cost when designing future optical networks.

12.5 Conclusion

We presented an adaptive routing and wavelength assignment protocol, the virtual circuit deflection (VCD) scheme, which is suitable for all-optical regular networks with wavelength conversion and outperforms oblivious routing schemes in the hypercube and torus topologies. We demonstrated that for the topologies considered, the performance of a system using the adaptive VCD scheme, with only one alternate link per hop, approaches that of a system using oblivious routing with twice as many wavelengths per link. We also presented performance results for the VCD protocol, focusing mainly on the adaptivity VCD can exhibit and noticed that when link capacity is large (i.e., more channels can simultaneously use the link and therefore VCD adaptivity is enhanced), the efficiency of the VCD protocol increases significantly. We also examined the effect of wavelength conversion in network performance, considering the cases of no-wavelength conversion, of limited-wavelength conversion and full-wavelength conversion and observed that limited conversion to only one or two adjacent wavelengths can provide a considerable fraction of the improvement that full-wavelength conversion provides over no-wavelength conversion. These results clearly emphasize the need for network designers to investigate the tradeoffs between wavelength conversion, routing flexibility, and hardware cost when designing future optical networks.

References

1. B. Mukherjee, WDM optical communications networks: progress and challenges, IEEE J. Select. Areas Commun., vol. 18, Oct. 2000, pp. 1810–1823.
2. H. Zang et al., Dynamic lightpath establishment in wavelength-routed WDM networks, IEEE Commun.s Mag., Sep. 2001, pp.100–108.
3. R. Ramaswami and A. Segall, Distributed network control for optical networks, IEEE/ACM Trans. Networking, vol. 5, no. 6, Dec. 1997, pp. 936–943.
4. H. Zang et al., Connection management for wavelength-routed WDM networks, Proc. IEEE GLOBECOM '99, Dec. 1999, pp. 1428–1432.
5. H. Zang, J. Jue, and B. Mukherjee, A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks, Opt. Networking Mag., vol. 1, no. 1, Jan. 2000, pp. 47–60.

6. J. Lang, V. Sharma, and E. Varvarigos, An analysis of oblivious and adaptive routing in optical networks with wavelength translation, IEEE/ACM Trans. Networking, vol. 9, no. 4, Aug. 2001, pp. 503–517.
7. R. Barry and P. Humblet, Models of blocking probability in all-optical networks with and without wavelength changers, IEEE J. Selected Areas Commun., vol. 14, June 1996, pp. 858–867.
8. M. Kovacevic and A. Acampora, Benefits of wavelength translation in all-optical clear-channel networks, IEEE J. Selected Areas Commun., vol. 14, June 1996, pp. 868–880.
9. V. Sharma and E.A. Varvarigos, An analysis of limited wavelength translation in regular all-optical WDM networks, Proc. IEEE INFOCOM '98, vol. 2, Mar.–Apr. 1998, pp. 893–901.
10. H. Harai, M. Murata, and H. Miyahara, Performance analysis of wavelength assignment policies in all-optical networks with limited-range wavelength conversion, *IEEE J. Selected Areas Commun.,* vol. 16, Sep. 1998, pp. 1051–1060.
11. B. Ramamurthy and B. Mukherjee, Wavelength conversion in WDM networking, *IEEE J. Selected Areas Commun.,* vol. 16, Sep. 1998, pp. 1061–1073.
12. K.N. Sivarajan and R. Ramaswami, Lightwave networks based on de Bruijn graphs, *IEEE/ACM Trans. Networking,* vol. 2, Feb. 1994, pp. 70–79.
13. Y. Zhu, G. Rouskas, and H. Perros, Blocking in wavelength routing networks. Part 1: The single path case, *Proc., IEEE INFOCOM '99,* vol. 1, Mar. 1999, pp. 321–328.
14. A. Karasan and E. Ayanoglu, Effects of wavelength routing and wavelength selection algorithms on wavelength conversion gain in WDM optical networks, *IEEE/ACM Trans. Networking,* vol. 6, Apr. 1998, pp. 186–196.
15. A. Mokhtar and M. Azizoglu, A new analysis of wavelength translation in regular WDM networks, *IEEE/ACM Trans. Networking,* vol. 6, Apr. 1998, pp. 197–206.
16. H. Harai, M. Murata, and H. Miyahara, Performance of alternate routing methods in all-optical switching networks, *Proc. IEEE INFOCOM '97,* vol. 2, Apr. 1997, pp. 516–524.
17. E.A. Varvarigos and J. Lang, A virtual circuit deflection protocol, *IEEE/ACM Trans. Networking,* vol. 7, no. 3, June 1999, pp. 335–349.
18. R. Koch, Increasing the size of a network by a constant factor can increase performance by more than a constant factor, *Symp. Foundations of Computer Science,* vol. 18, Oct. 1988, pp. 221–230.
19. S. Subramaniam, M. Azizoglu, and A.K. Somani, All-optical networks with sparse wavelength conversion, *IEEE/ACM Trans. Networking,* vol. 4, no. 4, Aug. 1996, pp. 544–557.
20. J. Yates, J. Lacey, D. Everitt, and M. Summerfield, Limited-range wavelength translation in all-optical networks, *Proc. IEEE INFOCOM '96*, vol. 3, Mar. 1996, pp. 954–961.
21. R. Ramaswami and G.H. Sasaki, Multi-wavelength optical networks with limited wavelength conversion, *Proc. INFOCOM '97,* Kobe, Japan, Apr. 1997.
22. N. Wauters and P. Demeester, Wavelength translation in optical multi-wavelength multi-fiber transport networks, *Int. J. Optoelectronics*, vol. 11, no. 1, Jan.–Feb. 1997, pp. 53–70.

Author: Ref. 12 is not cited in text.

Author: Ref. 13 is not cited in text.

23. N.F. Maxemchuk, Problems arising from deflection routing: live-lock, lock-out, congestion and message reassembly, in *High-Capacity Local and Metropolitan Area Networks,* G. Pujolle, Ed., Berlin, Germany: Springer-Verlag, June 1990.
24. F. Borgonovo, L. Fratta, and J. Bannister, Unslotted deflection routing in all-optical networks, *Proc. IEEE GLOBECOM '93*, vol. 1, pp. 119–125.
25. Z. Haas and D.R. Cheriton, Blazenet: a packet-switched wide-area network with photonic data path, *IEEE Trans. Commun.,* vol. 38, pp. 818–829, June 1990.
26. A.G. Greenberg and J. Goodman, Sharp approximate models of adaptive routing in mesh networks, *Proc. Teletraffic Analysis and Computer Performance Evaluation,* June 1986, pp. 255–270.
27. A. G. Greenberg and B. Hajek, Deflection routing in hypercube networks, *IEEE Trans. Commun.,* vol. 40, June 1992, pp. 1070–1081.
28. J. Bannister, F. Borgonovo, L. Fratta, and M. Gerla, A versatile model for predicting the performance of deflection-routing networks, *Perform. Eval.,* vol. 16, nos. 1–3, Nov. 1992, pp. 201–222.
29. F. Borgonovo, L. Fratta, and J. Bannister, On the design of optical deflection-routing networks, *Proc. IEEE INFOCOM '94,* vol. 1, pp. 120–129.
30. N.F. Maxemchuk, Comparison of deflection and store-and-forward techniques in the Manhattan Street and Shuffle-Exchange networks, *Proc. IEEE INFOCOM '89*, vol. 3, pp. 800–809.
31. J.T. Brassil, Deflection routing in certain regular networks, Ph.D. dissertation, University of California at San Diego, 1991.
32. A. Krishna and B. Hajek, Performance of shuffle-like switching networks with deflection, *Proc. IEEE INFOCOM '90*, vol. 2, pp. 473–480.
33. E.A. Varvarigos and D.P. Bertsekas, Performance of hypercube routing schemes with or without buffering, *IEEE/ACM Trans. Networking,* vol. 2, June 1994, pp. 299–311.
34. V. Sharma, Efficient communication protocols and performance analysis for gigabit networks, Ph.D. dissertation, Department ECE, University of California, Santa Barbara, Aug. 1997.

Author: Please spell out.

35. D. Bertsekas and R. Gallager, *Data Networks,* Englewood Cliffs, NJ: Prentice-Hall, 1992.
36. A. Butner and D. Skirmont, Architecture and design of a 40 gigabit per second ATM switch, *Proc. ICCD '95,* pp. 352–357.
37. E.A. Varvarigos and V. Sharma, The ready-to-go virtual circuit protocol: a loss-free protocol for multigigabit networks using FIFO buffers, *IEEE/ACM Trans. Networking,* vol. 5, Oct. 1997, pp. 705–718.