Optimal Virtual Topologies for One-To-Many Communication in WDM Paths and Rings

Jeff R. K. Hartline, Ran Libeskind-Hadas, *Member, IEEE*, Kurt M. Dresner, Ethan W. Drucker, and Katrina J. Ray

***Abstract*—In this paper we examine the problem of constructing optimal virtual topologies for one-to-many communication in optical networks employing wavelength-division multiplexing. A virtual topology is a collection of optical lightpaths embedded in a physical topology. A packet sent from the source node travels over one or more lightpaths en route to its destination. Within a lightpath, transmission is entirely optical. At the terminus of a lightpath the data is converted into the electronic domain where it may be retransmitted on another lightpath toward its destination. Since the conversion of the packet from the optical to the electronic domain introduces delays and uses limited physical resources, one important objective is to find virtual topologies which minimize either the maximum or average number of lightpaths used from the source to all destination nodes. Although this problem is NP-complete in general, we show that minimizing the maximum or average number of lightpaths in path and ring topologies can be solved optimally by efficient algorithms.**

I. Introduction

WAVELENGTH-DIVISION MULTIPLEXING (WDM) has emerged as a key optical networking technology for realizing low cost, high bandwidth, and scalable data services. Each fiber optical physical link in a WDM network is partitioned into multiple data channels each of which operates on a separate wavelength. Thus, WDM permits use of enormous fiber bandwidth by providing data channels whose individual bandwidths more closely match those of the electronic devices at their endpoints [13]. As WDM technology matures, it is likely to be widely used in systems ranging from local and metropolitan area networks to the backbone of the Next Generation Internet.

Manuscript received March 29, 2002; revised March 3, 2003; approved by IEEE/ACM Transactions on Networking Editor C. Qiao. This work was supported by the National Science Foundation under Grant CCR-9900491 to Harvey Mudd College.

J. R. K. Hartline was with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA. He is currently with the Department of Computer Science, Cornell University, Ithaca, NY 14853 USA (e-mail: jhartlin@cs.cornell.edu).

R. Libeskind-Hadas is with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA (e-mail: hadas@cs.hmc.edu).

K. M. Dresner was with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA. He is currently with the Department of Computer Science, University of Texas, Austin, TX 78712 USA (e-mail: kdresner@cs.hmc.edu).

E. W. Drucker was with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA. He is currently with the Department of Computer Science, University of California, Los Angeles, CA 90095 USA (e-mail: ethan@cs.ucla.edu).

K. J. Ray was with the Department of Computer Science, Harvey Mudd College, Claremont, CA 91711 USA. She is currently with the Department of Computer Science, New Mexico State University, Las Cruces, NM 88003 USA (e-mail: kray@cs.nmsu.edu).

Digital Object Identifier 10.1109/TNET.2004.826283

In many applications, a single source node in a network is required to send data to a number of destination nodes. The data sent to the destination nodes may be identical (multicast communication) or may be personalized. Ideally, each message is transmitted from the source to a destination without any optical-to-electronic conversion within the network. Such *all-optical* communication can be realized by using a single wavelength to establish a connection to each destination, but such connections require many dedicated optical paths which may, in general, be difficult or impossible to find [11]. Alternatively, all-optical wavelength converters may be used to convert from one wavelength to another within the network but such converters are likely to be prohibitively expensive for most applications in the foreseeable future [13]. Moreover, in all-optical communication a path is typically dedicated for communication from the source to a specific destination, potentially under-utilizing the bandwidth of the channels on that path.

A second approach is to embed a set of *lightpaths* in the network. A lightpath is a path comprising channels on a single wavelength. A packet may travel over multiple lightpaths from the source to a destination node. Within a lightpath, transmission is entirely optical. At the terminus of a lightpath the data is converted into electronic form and is delivered to the local node if it has reached its destination and is otherwise retransmitted on another lightpath toward its destination. Intermediate nodes on a lightpath simply allow the packet to pass through optically, but do not have access to the packets themselves. Thus, a single lightpath from node x to z passing through intermediate node y is different from two lightpaths, one from x to y and one from y to z, even if these two lightpaths use the same wavelength. The collection of lightpaths is called the *virtual topology*.

In a virtual topology, the *hop distance*[1] from node s to node d is defined to be the minimum number of lightpaths over all paths from s to d. Each hop in a path incurs latency due to conversion of the packet from the optical to the electronic domain, application of the routing function, queueing, and retransmission on the next lightpath. Thus, minimizing the number of hops is a performance metric of interest in networks in which hop latency dominates channel propagation latency [3], [14]. It has also been observed that under certain assumptions, minimizing the average hop distance is meaningful in wide area networks as well [2]. In addition to minimizing the maximum or average hop distance, other optimization criteria that are typically considered are minimizing congestion, minimizing use of certain network resources (transceivers, add/drop multiplexers, wavelength converters, among others), and minimizing propagation delay.

The virtual topology problem has been widely studied in the literature for a variety of network models and optimization

[1]This is also sometimes called *virtual hop distance*.

1063-6692/04$20.00 © 2004 IEEE

criteria. Recently, several aspects of virtual topology design have been considered for one-to-many communication in switch-based networks. Due to the apparent intractability of these problems, most existing solutions use either integer linear programming formulations or heuristics. The integer linear programming formulations allow optimal solutions, although they require exponential time in the worst case. In contrast, heuristics are generally fast but provide no guarantees on the optimality of the solutions. For example, Sahasrabuddhe and Mukherjee [10] use a mixed-integer linear programming formulation to minimize the average hop distance or the average number of transceivers used in a network. Zhang *et al.* [15] proposed and evaluated several heuristics for sparse-splitting networks in which only a subset of the nodes support multicast in hardware. Liang and Shen [8] examined the problem of minimizing the total cost of a multicast tree in which a cost is associated with each wavelength on each link and with converting from one wavelength to another. They showed that finding a virtual topology of minimum total cost in this model is NP-complete but that a polynomial-time approximation algorithm exists for this problem. Examples of other related results can be found in [1], [7], and [10]. Thaker and Rouskas [12] provide an excellent survey of results on multicasting in broadcast star networks. Finally, we note that the problem of finding optimal virtual topologies has also been studied in other types of networks. For example, related problems arise in ATM networks [4], [5].

Although the virtual topology problem in switch-based networks is computationally intractable in many cases, in this paper we demonstrate that the problem can be solved optimally in polynomial time for paths and rings under the objective of minimizing the average or maximum hop distance. These results are of potential practical and theoretical significance. In practice, rings are among the most prevalent topologies used in WDM networks. Moreover, the techniques and results described here may be useful in obtaining similar results for other related topologies.

The remainder of this paper is organized as follows. In Section II we describe the problem formally and give some preliminary results which will be used in the remainder of the paper. In Section III we give the fundamental result for paths and rings. In Section IV we generalize this result to the case of minimizing the weighted average hop distance. We conclude in Section V.

II. Problem Formulation and Preliminaries

We model a network by a directed graph $G = (V, E)$. We assume that the graph is symmetric so that if $(u, v) \in E$ then $(v, u) \in E$. We henceforth use the terms "vertex" and "node" interchangeably and the terms "directed edge" and "physical link" interchangeably.

Let $s \in V$ be the source vertex and $D \subseteq V - s$ be a set of destination vertices. The pair (s, D) is henceforth called a *communication group*. Let w be the number of wavelengths on each physical link which may be used to build the virtual topology for communication group (s, D). (Note that this may be the actual number of wavelengths present or may be a bound established in order to reserve wavelengths for other virtual topologies.) A *w-wavelength virtual topology* for the group (s, D) is a pair $T = (P, f)$ where P is a set of directed paths in G called *lightpaths* and $f : P \rightarrow \{1, \ldots, w\}$ is an assignment of lightpaths to wavelengths such that:

- Any two lightpaths in P which share a directed edge are assigned distinct wavelengths. That is, if $p, p' \in P$ and $\exists e \in E$ such that $e \in p \cap p'$ then $f(p) \neq f(p')$.
- For each $d \in D$ there exists a set of lightpaths $\{p_1, \ldots, p_k\} \subseteq P$ such that p_1 originates at s, p_k terminates at d, and for $1 \leq i \leq k-1$, the end vertex of p_i is the starting vertex of p_{i+1}.

Henceforth, we use the term *virtual path* to indicate a path from the source to a destination vertex comprising one or more lightpaths.

For a fixed communication group (s, D) the *hop distance* from s to d in virtual topology T, denoted $h_T(d)$, is defined to be the minimum number of lightpaths over all virtual paths from s to d. *The maximum hop distance* with respect to T is $\mathcal{M}(T) = \max_{d \in D}\{h_T(d)\}$ and the *average hop distance* with respect to T is $\mathcal{A}(T) = (\sum_{d \in D} h_T(d)/|D|)$.

Our objective is to find virtual topologies which are optimal with respect to metrics $\mathcal{M}$ or $\mathcal{A}$. The following theorem states that the optimal virtual topology problem is, in general, computationally intractable for both of these metrics.

Theorem 1: The problem of minimizing the maximum hop distance or average hop distance in an arbitrary symmetric directed graph is NP-complete.

The proof of this theorem is somewhat involved and is omitted in the interest of brevity. The interested reader is referred to [6].

Although the general problem is NP-complete, in the remainder of the paper we show that optimal virtual topologies can be found in polynomial time for symmetric paths and rings. We note that related problems were studied by Gerstel *et al.* in the context of ATM networks [4], [5]. In that model the path is unidirectional and thus all virtual paths are oriented in the same direction. This is not required in the WDM model, which makes the algorithms and analysis substantially more complicated.

Underlying our algorithms are two lemmas which are stated and proved below.

Lemma 1: Let $G = (V, E)$ be a symmetric graph and (s, D) a communication group in the graph. Let T' be any virtual topology for (s, D). There exists a virtual topology T such that for each $d \in D$ exactly one lightpath terminates at d and $h_T(d) \leq h_{T'}(d)$.

Proof: Let T' be an arbitrary virtual topology for (s, D) and let $d \in D$ be a vertex such that two or more lightpaths terminate at d. Let π be a virtual path in T' from s to d comprising the minimum number of lightpaths. The lightpath in π terminating at d is preserved and the remaining lightpaths terminating at d are eliminated. This transformation does not increase the hop distance to any vertex in D. This process is repeated until exactly one lightpath terminates at each $d \in D$. The resulting virtual topology is T. □

Next, we define the concept of crossing-free virtual topologies in paths. Let P be a symmetric path of n vertices $v_1, \ldots, v_n$. Let s be one of the vertices in the path and let D be a subset of vertices in the path excluding s. Let T' be an arbitrary virtual topology for communication group (s, D). Let indices i, j, k, ℓ be such that $i < j < k < \ell$. Let p and p' be two lightpaths in T' such that p originates at v_i and terminates at v_k, or vice versa, and p' originates at v_j and terminates at v_ℓ, or vice versa. In this case, lightpaths p and p' are said to *cross*. The four possible ways in which p and p' can cross are illustrated in Fig. 1. A virtual topology for a path in which no lightpaths cross is said to be *crossing-free*.

Lemma 2: Let $P = v_1, \ldots, v_n$ be a symmetric path and (s, D) a communication group in the path. Let T' be an arbitrary virtual topology for (s, D). There exists a virtual topology T for (s, D) such that T is crossing-free and $h_T(d) \leq h_{T'}(d)$ for every $d \in D$.

Proof: If T' is crossing-free then the lemma is vacuously true. Assume, therefore, that T' has at least one pair of crossing lightpaths. By Lemma 1, without loss of generality we may assume that there is a unique virtual path in T' from s to each $d \in D$.

We first show that for each of the four possible ways in which a pair of lightpaths can cross (as indicated in Fig. 1), the two crossing lightpaths can be replaced by a new pair of lightpaths which do not cross, do not introduce new crossings, and do not increase the hop distance from s to any vertex in the network.

Let p and p' denote a pair of crossing lightpaths and let λ and λ' denote the wavelengths assigned to these lightpaths, respectively. Note that λ and λ' may or may not be equal, depending on the relative orientations of p and p'.

First, consider the case in Fig. 1(a). In this case the wavelength λ assigned to p is distinct from the wavelength λ' assigned to p'. If $h_{T'}(v_i) < h_{T'}(v_j)$ then replace p' with a lightpath from v_k to v_ℓ assigned to wavelength λ'. The new hop distances to v_i, v_j, and v_k are unchanged and the new hop distance to v_ℓ is $h_{T'}(v_i)+2 \leq h_{T'}(v_j)+1 = h_{T'}(v_\ell)$. No other hop distances are affected and no new crossings are introduced. If $h_{T'}(v_i) \geq h_{T'}(v_j)$ then remove p from T' and add a lightpath from v_j to v_k assigned to wavelength λ. The hop distance to v_k decreases or remains unchanged and all other hop distances remain unchanged. The case in Fig. 1(b) is analogous.

Next, consider the case in Fig. 1(c). If $h_{T'}(v_i) < h_{T'}(v_\ell)$ then replace p' with a virtual path from v_k to v_j assigned to wavelength λ'. Note that all hop distances remain unchanged with the possible exception of the distance to v_j which is now $h_{T'}(v_i)+2 \leq h_{T'}(v_\ell)+1 = h_{T'}(v_j)$. The case that $h_{T'}(v_i) > h_{T'}(v_\ell)$ is analogous. In both cases, no new crossings are introduced. If $h_{T'}(v_i) = h_{T'}(v_\ell)$ then replace p and p' with virtual paths from v_i to v_j and v_ℓ to v_k assigned to wavelengths λ and λ', respectively. All hop distances remain unchanged and no new crossings are introduced.

Finally, consider the case in Fig. 1(d). If $h_{T'}(v_j) < h_{T'}(v_k)$ then replace p with a virtual path from v_j to v_i assigned to wavelength λ. The case for $h_{T'}(v_j) > h_{T'}(v_k)$ is analogous. In both cases no hop distances are increased and no new crossings are introduced. If $h_{T'}(v_j) = h_{T'}(v_k)$ then replace p with a virtual path from v_j to v_i assigned to wavelength λ and replace p' with a virtual path from v_k to v_ℓ assigned to wavelength λ'. All hop distances are unchanged and no new crossings are introduced.

Note that each of the four transformations described above reduces the total number of crossings. Thus, given a virtual topology T' with c crossings the transformations can be applied at most c times until no crossings remain, resulting in a new crossing-free virtual topology T such that $h_T(d) \leq h_{T'}(d)$ for all $d \in D$. □

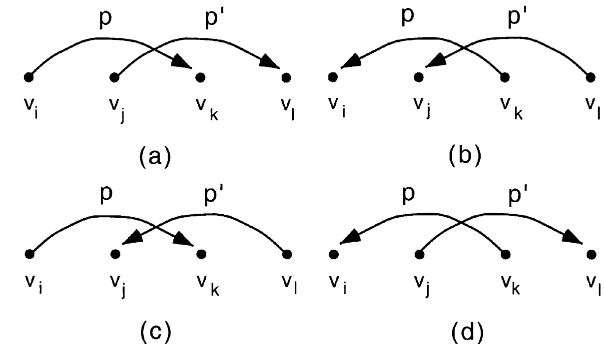

Fig. 1. The four possible ways in which two paths p and p' can cross.

Since our objectives are to minimize the maximum or average hop distance in a virtual topology, Lemma 2 implies that we may henceforth restrict our attention to crossing-free virtual topologies.

In the next section we use these results to show that the optimal topology problem, with respect to $\mathcal{M}$ or $\mathcal{A}$, can be solved efficiently in paths and rings. Moreover, we show how to construct virtual topologies that simultaneously minimize both the maximum and the average hop distance in these topologies. In Section IV we generalize the problem formulation by associating a positive real number weight with each destination node in D. We give polynomial-time dynamic programming algorithms for minimizing the average weighted hop distance or the maximum weighted hop distance in paths and rings.

III. Unweighted Paths and Rings

In this section we characterize the relationship between the number of nodes, the number of wavelengths, and the maximum and average hop distance in a virtual topology for paths and rings. In addition, we show how to construct optimal virtual topologies with respect to minimizing both the maximum and the average hop distance in these topologies. We begin by considering symmetric directed paths. By Lemma 2 we assume, without loss of generality, that all virtual topologies are crossing-free. Note that if the source vertex s is an interior vertex in a path, we may assume without loss of generality that an optimal virtual topology contains no lightpath which originates at a vertex u to the right (respectively, left) of s and terminates at a vertex v to the left (respectively, right) of s; any such topology could be trivially transformed into a virtual topology in which the terminus v of the lightpath receives the message in fewer hops directly from s. Thus, if s is an interior vertex in a path, we may simply partition the path into two paths, one containing s and all vertices to its right and one containing s and all vertices to its left. The two virtual topology problems can then be solved as two independent problems in which the source is the leftmost and rightmost vertex, respectively. Since the cases that the source is the left or right endpoint of the path are symmetric, we henceforth assume without loss of generality that the source s is the leftmost vertex in the path.

Given positive integers h and w, we first find the longest possible path for which there exists a virtual topology such that at most w wavelengths are used per directed edge, each vertex in the path can be reached within at most h hops, and for each hop value i between 0 and h the number of vertices within i hops of the source is maximized. More formally, we find a virtual topology $\mathcal{T}(h, w)$ in a symmetric directed path with the following properties:

1) $\mathcal{T}(h, w)$ uses at most w wavelengths per directed edge.
2) The maximum hop distance in $\mathcal{T}(h, w)$ is h.
3) Let T' be any crossing-free virtual topology in a path using up to w wavelengths per directed edge. For each $0 \le i \le h$, the number of vertices in T' with hop distance i is no larger than the number of vertices in $\mathcal{T}(h, w)$ with hop distance i.

We then find a closed form for the number of nodes in $\mathcal{T}(h, w)$. Finally, we use these results to construct optimal virtual topologies with respect to $\mathcal{M}$ and $\mathcal{A}$ in paths and then rings.

Although we assume that each edge has the same number, w, of wavelengths available, our subsequent analysis is simplified by defining a topology $T(h, w_{\text{out}}, w_{\text{in}})$ which uses at most w_{out} wavelengths on edges oriented out away from the source and at most w_{in} wavelengths on edges oriented in toward the source. Let $\overline{T}(h, w_{\text{out}}, w_{\text{in}})$ denote the reflection of $T(h, w_{\text{out}}, w_{\text{in}})$ with respect to s.

Definition 1: $T(h, w_{\text{out}}, w_{\text{in}})$ is constructed recursively as follows:

1) If $h = 0$ or $w_{\text{out}} = 0$ then $T(h, w_{\text{out}}, w_{\text{in}})$ is a single vertex (the source vertex, s).
2) If $h > 0$ and $w_{\text{out}} > 0$ then $T(h, w_{\text{out}}, w_{\text{in}})$ comprises a lightpath from s to a vertex v on wavelength w_{out}. Vertex v is the source of $T(h-1, w_{\text{out}}, w_{\text{in}})$ to its right and $\overline{T}(h-1, w_{\text{in}}, w_{\text{out}}-1)$ to its left. Vertex s is also the source of $T(h, w_{\text{out}}-1, w_{\text{in}})$ to its right.

This construction is shown Fig. 2.[2] For example, $T(2,2,2)$ is shown in Fig. 3 where lightpaths using one of the two wavelengths are shown using solid edges and lightpaths using the second wavelength are shown using dashed edges. Note that $T(h, w_{\text{out}}, w_{\text{in}})$ is crossing-free by construction.

Theorem 2: Let $T'(w_{\text{out}}, w_{\text{in}})$ be an arbitrary crossing-free virtual topology in a path using at most w_{out} wavelengths on each link oriented out away from s and at most w_{in} wavelengths on each link oriented in toward s. For each i, $0 \le i \le h$, the number of vertices in $T(h, w_{\text{out}}, w_{\text{in}})$ with hop distance i is at least as large as the number of vertices in $T'(w_{\text{out}}, w_{\text{in}})$ with hop distance i.

Proof: The proof is by strong induction on $h + w_{\text{out}} + w_{\text{in}}$.

Basis: When $h + w_{\text{out}} + w_{\text{in}} = 0$, $h = w_{\text{out}} = w_{\text{in}} = 0$. In this case, no vertices are reachable from the source and the only possible virtual topology is a single vertex. By Definition 1, $T(0,0,0)$ comprises a single vertex.

[2]Note that $T(h, w_{\text{out}}, w_{\text{in}})$ is well-defined although there is no base case given for $w_{\text{in}} = 0$. If $h > 0$, $w_{\text{out}} > 0$, and $w_{\text{in}} = 0$, then none of $T(h, w_{\text{out}}-1, w_{\text{in}})$, $\overline{T}(h-1, w_{\text{in}}, w_{\text{out}}-1)$, and $T(h-1, w_{\text{out}}, w_{\text{in}})$ contain negative arguments and thus each is well-defined.

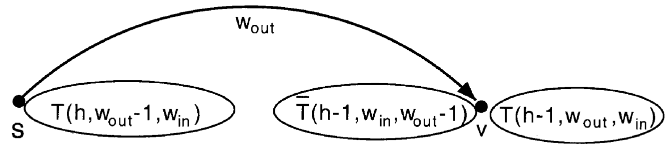

Fig. 2. The construction of $T(h, w_{\text{in}}, w_{\text{out}})$ in Definition 1.

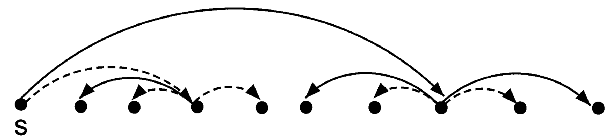

Fig. 3. $T(2,2,2)$ with one wavelength indicated with solid edges and the other wavelength indicated with dashed edges.

Induction Step: Assume that the claim is true for $0 \le h + w_{\text{out}} + w_{\text{in}} < n$ and consider the case that $n = h + w_{\text{out}} + w_{\text{in}}$. If $h = 0$ then by construction $T(h, w_{\text{out}}, w_{\text{in}})$ comprises only the source vertex and the claim is trivially true. If $w_{\text{out}} = 0$ then both $T'(w_{\text{out}}, w_{\text{in}})$ and $T(h, w_{\text{out}}, w_{\text{in}})$ comprise only the source vertex and the claim is also true. Thus, we may assume that $h > 0$ and $w_{\text{out}} > 0$.

Consider crossing-free virtual topologies $T'(w_{\text{out}}, w_{\text{in}})$ and $T(h, w_{\text{out}}, w_{\text{in}})$. In $T'(w_{\text{out}}, w_{\text{in}})$, let u be the rightmost vertex connected to s by a single lightpath and let p denote this lightpath. Without loss of generality, p is assigned wavelength w_{out}. Without loss of generality, all destination vertices to the right of u are reached using lightpath p. If not, there exists a destination vertex y to the right of u which is reached via a vertex at x to the left of u. The path from x to y comprises one or more lightpaths. If the path comprises one lightpath, then this lightpath crosses p, contradicting the assumption that T' is crossing-free. Otherwise, the path from x to y comprises two or more lightpaths. Since T' is crossing-free, one lightpath on the path from x to y must terminate at u and one lightpath must originate at u. In this case, y can be reached using fewer hops by first using lightpath p and then continuing from u as in the original path in T' from x to y.

Let v be the rightmost vertex in $T(h, w_{\text{out}}, w_{\text{in}})$ connected to s by a single lightpath, as in the construction in Fig. 2. By construction, the vertices to the right of v in $T(h, w_{\text{out}}, w_{\text{in}})$ are the destination vertices of $T(h-1, w_{\text{out}}, w_{\text{in}})$. By the induction hypothesis, for each i, $0 \le i \le h-1$, the number of vertices in $T(h-1, w_{\text{out}}, w_{\text{in}})$ to the right of v with hop distance i from v is at least as large as the number of vertices to the right of u in $T'(w_{\text{out}}, w_{\text{in}})$ with hop distance i from u.

Next, consider the set of destination vertices in $T'(w_{\text{out}}, w_{\text{in}})$ to the right of s but to the left of u. Let D_u denote the set of these destination vertices that are reached using the lightpath from s to u and let D_s denote the remaining vertices. By Lemma 2, we may assume that the vertices in D_u are strictly to the right of the vertices in D_s. Since wavelength w_{out} is used for the lightpath from s to u, only wavelengths $1, \ldots, w_{\text{out}}-1$ are available in the out direction while all wavelengths $1, \ldots, w_{\text{in}}$ are available in the in direction between vertices s and u. Thus, in $T'(w_{\text{out}}, w_{\text{in}})$, the destination vertices in D_s are reached from s using up to $w_{\text{out}}-1$ wavelengths in the out direction and

w_{in} wavelengths in the in direction. Since the edges oriented in toward u and out those oriented out away from u are those oriented out away from s and in toward s, respectively, the destination vertices in D_u are reached from u using up to w_{in} wavelengths on edges oriented away from u and up to $w_{\text{out}}-1$ wavelengths on edges oriented toward u. By construction, the vertices in $T(h, w_{\text{out}}, w_{\text{in}})$ to the right of s and to the left of v are the destination vertices of $T(h, w_{\text{out}}-1, w_{\text{in}})$ and $\overline{T}(h-1, w_{\text{in}}, w_{\text{out}}-1)$. By the induction hypothesis, for each i, $0 \leq i \leq h$, the number of vertices in $T(h, w_{\text{out}}-1, w_{\text{in}})$ with hop distance i from s is at least as large as the number of vertices in D_s with the same hop distance. Similarly, by the induction hypothesis, for each i, $0 \leq i \leq h-1$, the number of vertices in $T(h, w_{\text{in}}, w_{\text{out}}-1)$ with hop distance i from v is at least as large as the number of vertices in D_u with the same hop distance. □

Since we are interested in the case that the number of wavelengths available on each edge is the same, regardless of the orientation of the edge, we now introduce the following definition:

Definition 2: $\mathcal{T}(h, w)$ is defined to be $T(h, w, w)$.

Corollary 1: Let $\mathcal{T}'(w)$ be an arbitrary crossing-free virtual topology in a path which uses at most w wavelengths on each link. For each i, $0 \leq i \leq h$, the number of vertices in $\mathcal{T}(h, w)$ with hop distance i is at least as large as the number of vertices in $\mathcal{T}'(w)$ with hop distance i.

Proof: Follows from Theorem 2 and Definition 2. □

Let $\mathcal{N}(h, w)$ denote the total number of vertices in $\mathcal{T}(h, w)$. We next derive a closed form for $\mathcal{N}(h, w)$. Although the inductive definition of $T(h, w_{\text{out}}, w_{\text{in}})$ induces a recurrence relation, this recurrence is evidently difficult to solve directly. We therefore employ a combinatorial argument to find an expression for $\mathcal{N}(h, w)$. To simplify the combinatorial argument let $\hat{\mathcal{T}}(h, w)$ be the superposition of $\mathcal{T}(h, w)$ and its reflection with respect to the source vertex, s. We count the total number of vertices in $\hat{\mathcal{T}}(h, w)$, which we denote $\hat{\mathcal{N}}(h, w)$. Note that $\mathcal{N}(h, w) = (1/2)(\hat{\mathcal{N}}(h, w) + 1)$ because the source vertex is in both $\hat{\mathcal{T}}(h, w)$ and its reflection.

We begin by defining the concept of a w-string. It is easily shown that there is a bijection between the w-strings of length at most h and the set of all vertices in $\hat{\mathcal{T}}(h, w)$. We then find a closed form for the number of w-strings of length at most h.

Definition 3: Let w be a nonnegative integer. Let $\Sigma_w = \{\overrightarrow{1}, \overrightarrow{2}, \ldots, \overrightarrow{w}, \overleftarrow{1}, \overleftarrow{2}, \ldots, \overleftarrow{w}\}$ if $w \geq 1$ and $\emptyset$ otherwise. A string in Σ_w^* is said to be a w-string if it satisfies the following two properties:

1) If $\overrightarrow{i}$ appears to the left of $\overrightarrow{j}$, or $\overleftarrow{i}$ appears to the left of $\overleftarrow{j}$, then $i \geq j$ and
2) All occurrences of a given symbol are contiguous.

For example, the following are 3-strings: $\overrightarrow{3}\,\overrightarrow{2}\,\overrightarrow{2}\,\overleftarrow{3}$, $\overleftarrow{2}\,\overrightarrow{3}\,\overrightarrow{2}\,\overleftarrow{1}$, $\overleftarrow{3}\,\overleftarrow{3}\,\overleftarrow{1}$, $\overleftarrow{2}\,\overleftarrow{1}\,\overrightarrow{3}$.

Theorem 3: There is a bijection between the w-strings of length at most h and the set of all vertices in $\hat{\mathcal{T}}(h, w)$.

Proof: For nonnegative integers w_{out} and w_{in} let $\Sigma_{w_{\text{out}}} = \{\overrightarrow{1}, \overrightarrow{2}, \ldots, \overrightarrow{w}_{\text{out}}\}$ if $w_{\text{out}} > 0$ and $\emptyset$ otherwise. Let $\Sigma_{w_{\text{in}}} = \{\overleftarrow{1}, \overleftarrow{2}, \ldots, \overleftarrow{w}_{\text{in}}\}$ if $w_{\text{in}} > 0$ and $\emptyset$ otherwise. A $(w_{\text{out}}, w_{\text{in}})$-*string* is defined to be a string over alphabet $(\Sigma_{w_{\text{out}}} \cup \Sigma_{w_{\text{in}}})^*$ such that if $\overrightarrow{i}$ appears to the left of $\overrightarrow{j}$, or $\overleftarrow{i}$ appears to the left of $\overleftarrow{j}$, then $i \geq j$ and all occurrences of a given symbol are contiguous. Let $S(h, w_{\text{out}}, w_{\text{in}})$ denote the set of all $(w_{\text{out}}, w_{\text{in}})$-strings of length at most h such that first symbol in the string is not an element of $\Sigma_{w_{\text{in}}}$. For each $(w_{\text{out}}, w_{\text{in}})$-string, s, let the *reflection of* s, denoted $\overline{s}$, be the string constructed by replacing each occurrence of $\overrightarrow{i}$ in s by $\overleftarrow{i}$ and replacing each occurrence of $\overleftarrow{i}$ in s by $\overrightarrow{i}$. For a set A of $(w_{\text{out}}, w_{\text{in}})$-strings, define the reflection of A, denoted $\overline{A}$, to be the set of reflections of all strings in A.

We claim that there is a bijection between the strings in $S(h, w_{\text{out}}, w_{\text{in}})$ and the set of all vertices in $T(h, w_{\text{out}}, w_{\text{in}})$. The theorem is then shown to follow immediately from the correctness of this claim.

The proof is by strong induction on $h + w_{\text{out}} + w_{\text{in}}$.

Basis: When $h + w_{\text{in}} + w_{\text{out}} = 0$, $h = w_{\text{in}} = w_{\text{out}} = 0$. In this case, there is exactly one $(w_{\text{in}} + w_{\text{out}})$-string, the empty string. There is also only one vertex, the source vertex, in $T(h, w_{\text{out}}, w_{\text{in}})$. Thus, a bijection exists between the two sets.

Induction Step: Assume that the claim is true for $0 \leq h + w_{\text{out}} + w_{\text{in}} < n$ and consider the case that $n = h + w_{\text{out}} + w_{\text{in}}$. If $h = 0$ or $w_{\text{out}} = 0$ then $S(h, w_{\text{out}}, w_{\text{in}})$ contains only the empty string. Otherwise, let s denote a string in $S(h, w_{\text{out}}, w_{\text{in}})$. The first symbol in s is either $\overrightarrow{w}_{\text{out}}$ or $\overrightarrow{i}$ where $i < w_{\text{out}}$. We consider these two cases separately.

Case 1) The first symbol in s is $\overrightarrow{w}_{\text{out}}$. Then the remainder of the string may begin with either a symbol in $\Sigma_{w_{\text{out}}}$ or $\Sigma_{w_{\text{in}}}$. In the former case, the remainder of the string is any string in $S(h-1, w_{\text{out}}, w_{\text{in}})$. In the latter case, the remainder of the string is any string in $\overline{S}(h-1, w_{\text{in}}, w_{\text{out}}-1)$. By the induction hypothesis, there is a bijection between $S(h-1, w_{\text{out}}, w_{\text{in}})$ and $T(h-1, w_{\text{out}}, w_{\text{in}})$ and also between $S(h-1, w_{\text{in}}, w_{\text{out}}-1)$ and $T(h-1, w_{\text{in}}, w_{\text{out}}-1)$ and therefore between $\overline{S}(h-1, w_{\text{in}}, w_{\text{out}}-1)$ and $\overline{T}(h-1, w_{\text{in}}, w_{\text{out}}-1)$.

Case 2) The first symbol in s is $\overrightarrow{i}$ where $i < w_{\text{out}}$. In this case, s is any string in $S(h, w_{\text{out}}-1, w_{\text{in}})$ and, by the induction hypothesis, there is a bijection between this set and $T(h, w_{\text{out}}-1, w_{\text{in}})$.

Therefore, there is a bijection between $S(h, w_{\text{out}}, w_{\text{in}})$ and $T(h, w_{\text{out}}, w_{\text{in}})$. This completes the inductive proof of the claim.

Now consider $S(h, w, w)$ and its reflection set $\overline{S}(h, w, w)$. From the above induction proof, there is a bijection between $T(h, w, w)$ and $S(h, w, w)$ and between $\overline{T}(h, w, w)$ (the reflection of $T(h, w, w)$) and $\overline{S}(h, w, w)$. Moreover, $\hat{\mathcal{T}}(h, w)$ is simply the superposition of $T(h, w, w)$ and $\overline{T}(h, w, w)$, which are disjoint with the exception of the source vertex. Also, $S(h, w, w)$ and $\overline{S}(h, w, w)$ are disjoint with the exception of the empty string and their union is exactly the set of w-strings of length at most h. □

We now give an expression for $\hat{\mathcal{N}}(h, w)$. We use the standard notation $\binom{n}{k}$ to denote the number of ways to choose k items from n without repetition and $\left(\!\binom{n}{k}\!\right)$ to denote the number of ways to choose k items from n with repetition allowed.[3]

[3]Recall that $\binom{n}{k} = (n!/k!(n-k)!)$ and $\left(\!\binom{n}{k}\!\right) = \binom{n+k-1}{k}$.

Theorem 4:

$$\hat{\mathcal{N}}(h,w) = \sum_{i=0}^{\min\{\lfloor \frac{h}{2} \rfloor, w\}} \binom{w}{i}^2 \left(\!\binom{2w+1}{h-2i}\!\right)$$

Proof: From Theorem 3 the total number of vertices in $\hat{\mathcal{T}}(h,w)$ is equal to the number of w-strings of length at most h. Let $\overrightarrow{\Sigma} = \{\overrightarrow{1}, \overrightarrow{2}, \ldots, \overrightarrow{w}\}$ and $\overleftarrow{\Sigma} = \{\overleftarrow{1}, \overleftarrow{2}, \ldots, \overleftarrow{w}\}$. Let a *right block* be defined as a maximal substring of symbols from $\overrightarrow{\Sigma}$ and a *left block* to be a maximal substring of symbols from $\overleftarrow{\Sigma}$. A *transition* is said to occur in a w-string whenever a right block is adjacent to a left block, in either order, or the string begins with a nonempty left block.

Since there can be at most one transition per symbol in a string, a w-string of length h has at most h transitions. Moreover, such a string can have no more than $2w$ transitions, one between each of the $2w$ distinct symbols. Thus, there are at most $\min\{h, 2w\}$ transitions. For each $0 \le i \le \min\{\lfloor (h/2) \rfloor, w\}$, let $t(i)$ be the set of w-strings of length at most h which have a total of either $2i$ or $2i+1$ transitions.

A string with either $2i$ or $2i+1$ transitions must be of the form $r_0 \ell_0 r_1 \ell_1 \ldots r_i \ell_i$ where each r_j and ℓ_j, $0 \le j \le i$, is a right block and left block, respectively. We note that r_0 or ℓ_i may be empty, but the remaining blocks may not be empty. If ℓ_i is empty there are $2i$ transitions and otherwise there are $2i+1$ transitions.

We count the number of strings in $t(i)$ as follows. We choose from $\overrightarrow{\Sigma}$ the first symbol in each of the nonempty right blocks $r_1, \ldots, r_i$. There are $\binom{w}{i}$ choices. Similarly, we choose from $\overleftarrow{\Sigma}$ the last symbol in each of the nonempty left blocks $\ell_0, \ldots, \ell_{i-1}$. There are again $\binom{w}{i}$ choices. Next, the remaining symbols are chosen. There are a total of $2w$ symbols to choose from and, in addition, we may choose the empty symbol in order to construct a string of length less than h. Thus, from $2w+1$ symbols we must choose $h-2i$ more symbols, with repetition allowed. Thus, there are $(\binom{2w+1}{h-2i})$ choices for the remaining symbols. By the definition of w-strings, the locations of these $h-2i$ symbols in the string are induced by the choice of the already selected symbols. Therefore, the total number of strings in $t(i)$ is

$$\binom{w}{i}^2 \left(\!\binom{2w+1}{h-2i}\!\right)$$

and the theorem follows. □

Corollary 2:

$$\mathcal{N}(h,w) = \frac{1}{2}\left[1 + \sum_{i=0}^{\min\{\lfloor \frac{h}{2} \rfloor, w\}} \binom{w}{i}^2 \left(\!\binom{2w+1}{h-2i}\!\right)\right]$$

Proof: Follows immediately from Theorem 4 and the definition of $\mathcal{N}(h,w)$. □

Theorem 5: Let $\pi = v_1, \ldots, v_n$ be a symmetric directed path where v_1 is the source vertex and all remaining vertices are destinations. Let w denote the number of wavelengths available on each directed edge. Let $h = \min\{x | \mathcal{N}(x,w) \ge n\}$. Let $\mathcal{T}$ denote the virtual topology constructed by removing excess vertices from $\mathcal{T}(h,w)$ in decreasing order of hop count distance from source v_1, until exactly n vertices remain. Virtual topology $\mathcal{T}$ is optimal for π with respect to both metrics $\mathcal{M}$ and $\mathcal{A}$.

Proof: By the definition of h in the Theorem, $\mathcal{N}(h-1,w) < n$. Also, from Corollary 1 it follows that $\mathcal{T}(h-1,w)$ and $\mathcal{T}(h,w)$ contain the same number of destination vertices with hop distance i, $0 \le i \le h-1$. Therefore, the only vertices removed from $\mathcal{T}(h,w)$ in the construction of $\mathcal{T}$ are vertices with hop distance h. Moreover, $\mathcal{T}$ must have at least one vertex with hop distance h and the number of vertices in $\mathcal{T}$ with hop distance i, $0 \le i < h$, is equal to the number of vertices with hop distance i in $\mathcal{T}(h,w)$.

Next, we show that $\mathcal{T}$ is optimal with respect to metric $\mathcal{M}$. Assume that $\mathcal{T}'$ is a virtual topology for π using at most w wavelengths per edge and $\mathcal{M}(\mathcal{T}') < \mathcal{M}(\mathcal{T}) = h$. Without loss of generality, $\mathcal{T}'$ is crossing-free. From the above observations and Corollary 1, it follows that for each i, $0 \le i < h$, $\mathcal{T}$ contains at least as many vertices at hop distance i as does $\mathcal{T}'$. However, $\mathcal{T}$ also has at least one vertex with hop distance h, and thus $\mathcal{T}$ contains a larger number of destination vertices than does $\mathcal{T}'$, contradicting the assumption that both $\mathcal{T}$ and $\mathcal{T}'$ are virtual topologies for path π with n vertices.

Finally, we show that $\mathcal{T}$ is optimal with respect to metric $\mathcal{A}$. Assume that $\mathcal{T}'$ is a virtual topology for π using at most w wavelengths per edge and $\mathcal{A}(\mathcal{T}') < \mathcal{A}(\mathcal{T})$. Without loss of generality, $\mathcal{T}'$ is crossing-free. Therefore, the sum of the hop distances with respect to $\mathcal{T}'$ is smaller than the sum of the hop distances with respect to $\mathcal{T}$. Let $d_1, \ldots, d_n$ and $d'_1, \ldots, d'_n$ be the hop distances in $\mathcal{T}$ and $\mathcal{T}'$, respectively, in nondecreasing order. By assumption, there exists i, $1 \le i \le n$, such that $d'_i < d_i$. Since $d_i \le h$ we have $d'_i < h$. Now, from the observations above and Corollary 1, we know that $\mathcal{T}$ has at least as many vertices with hop distance less than or equal to d'_i as does $\mathcal{T}'$. This contradicts the assumption that $d'_i < d_i$. □

From the recursive definition of $\mathcal{T}(h,w)$ it is easily verified that the optimal virtual topology $\mathcal{T}$ defined in Theorem 5 can also be constructed directly via a simple breadth-first procedure: First, w vertices within one hop of s are constructed and assigned lightpaths with wavelengths as defined in Definition 1. Next, all vertices within two hops of s are constructed and this process is repeated until n vertices are thus constructed. This breadth-first construction takes time $O(n)$.

Finally, we note that the extension of these results to rings is immediate. Consider a bidirectional ring with n nodes, one of which is the source and let w be the number of wavelengths available. By Lemma 2, we may restrict our attention to crossing-free virtual topologies. Thus, without loss of generality, a topology which is optimal with respect to metrics $\mathcal{M}$ and $\mathcal{A}$ can be partitioned into two disjoint virtual topologies, one which reaches a contiguous subset of nodes clockwise with respect to the source and one which reaches a subset of nodes counter-clockwise with respect to the source. In order to minimize the average or maximum hop distances, the ring is therefore partitioned into two sets of nodes of sizes $\lceil (n-1)/2 \rceil$ and $\lfloor (n-1)/2 \rfloor$, one clockwise and one counter-clockwise with respect to the source. Each of these is then treated as a separate path problem using the results in this section.

IV. Weighted Paths and Rings

In this section we consider the problem of minimizing the weighted average number of hops in a symmetric path under the assumption that each destination has a positive real weight associated with it. More precisely, let $v_0, v_1, \ldots, v_n$ be a bidirectional path with source node v_0 and let $c(v_i)$ be a positive real cost or weight associated with each destination node v_i, $1 \leq i \leq n$. For a given number of available wavelengths, w, our objective is to find a w-wavelength virtual topology T which minimizes the *weighted-average hop distance* defined as:

$$\mathcal{WA}(T) = \frac{1}{n} \sum_{i=1}^{n} h_T(v_i) c(v_i)$$

Note that the average hop distance is a special case of the weighted-average hop distance when all weights are equal to 1. The weighted-maximum hop distance could be defined analogously, but this is a straightforward analog of the weighted-average and is not described here in the interest of brevity. Henceforth, we use the term *optimal virtual topology* to mean optimal with respect to the weighted-average hop distance. For simplicity, we also refer to vertices in the path by their indices.

Let $\overrightarrow{W}(i,j,r,\ell)$ denote the minimum weighted-average hop distance for the case that the source is vertex i, $i \leq j$, the destination vertices are $i+1, \ldots, j$, and there are r and ℓ wavelengths available on each edge oriented to the right and left, respectively. Similarly, let $\overleftarrow{W}(i,j,r,\ell)$ denote the minimum weighted-average hop distance for the case that the source is vertex i, $i \geq j$, the destinations vertices are $i-1, \ldots, j$, and there are r and ℓ wavelengths available on each edge oriented to the right and left, respectively. Note that $\mathcal{WA}(T) = \overrightarrow{W}(0,n,w,w)$.

Theorem 6: If $i = j$ then $\overrightarrow{W}(i,j,r,\ell) = 0$ and $\overleftarrow{W}(i,j,r,\ell) = 0$. If $i < j$ and $r = 0$ then $\overrightarrow{W}(i,j,r,\ell) = \infty$ and if $i > j$ and $\ell = 0$ then $\overleftarrow{W}(i,j,r,\ell) = \infty$. Otherwise, for $i < j$,

$$\overrightarrow{W}(i,j,r,\ell) = \min_{i<s\leq j} \min_{i<t\leq s} \overrightarrow{W}(i,t-1,r-1,\ell) + \overleftarrow{W}(s,t,r-1,\ell) + \overrightarrow{W}(s,j,r,\ell) + \sum_{k=t}^{j} c(v_k)$$

and for $i > j$,

$$\overleftarrow{W}(i,j,r,\ell) = \min_{j\leq s<i} \min_{s\leq t<i} \overleftarrow{W}(i,t+1,r,\ell-1) + \overrightarrow{W}(s,t,r,\ell-1) + \overleftarrow{W}(s,j,r,\ell) + \sum_{k=j}^{t} c(v_k)$$

Proof: When $i = j$ there are no destination nodes and the weighted-average hop distance is 0. When $i < j$ and $r = 0$, no wavelengths are available to reach the nonzero number of destination nodes and thus $\overrightarrow{W}(i,j,r,\ell) = \infty$ since there exists no solution with finite average weight. Similarly, when $i > j$ and $\ell = 0$, $\overleftarrow{W}(i,j,r,\ell) = \infty$.

Next, consider the case that $i < j$ and $r > 0$. By Lemma 2, there exists an optimal virtual topology with source i and destinations $i+1, \ldots, j$ which is crossing-free. The rightmost vertex connected to i by a single lightpath has index s where $i < s \leq j$. Since the virtual topology is crossing-free, vertices $s+1, \ldots, j$ are reached via this lightpath. Thus, these vertices are reached from s using an optimal virtual topology with all r and ℓ wavelengths in the right and left directions available, respectively. However, each such node incurs an extra hop from i to s, and thus the contribution of these destinations to $\overrightarrow{W}(i,j,r,\ell)$ is $\overrightarrow{W}(s,j,r,\ell) + \sum_{k=s+1}^{j} c(v_k)$. Vertices $i+1, \ldots, s-1$ are either reached via s or not. Since the virtual topology is crossing-free, there exists an index t, $i < t \leq s$, such that vertices $t, \ldots, s-1$ are reached via s and vertices $i+1, \ldots, t-1$ are reached from i but not via s. Each of these two sets of destinations is reached using optimal topologies. Vertices $i+1, \ldots, t-1$ contribute $\overrightarrow{W}(i,t-1,r-1,\ell)$ to $\overrightarrow{W}(i,j,r,\ell)$ since one of the wavelengths in the right direction is used for the lightpath from i to s. Vertices $t, \ldots, s-1$ are visited from s using an optimal topology with source s and $r-1$ wavelengths available in the right direction and all ℓ wavelengths available in the left direction. However, the paths to these vertices incur an extra hop from i to s. Therefore, the contribution of these vertices to $\overrightarrow{W}(i,j,r,\ell)$ is $\overleftarrow{W}(s,t,r-1,\ell) + \sum_{k=t}^{s-1} c(v_k)$. Finally, s itself has hop distance 1 and thus contributes $c(v_s)$ to $\overrightarrow{W}(i,j,r,\ell)$. Therefore,

$$\overrightarrow{W}(i,j,r,\ell) = \min_{i<s\leq j} \min_{i<t\leq s} \overrightarrow{W}(i,t-1,r-1,\ell) + \overleftarrow{W}(s,t,r-1,\ell) + \overrightarrow{W}(s,j,r,\ell) + \sum_{k=t}^{j} c(v_k)$$

The case that $i > j$ and $r > 0$ is analogous. □

We can now use a dynamic programming algorithm to compute $\overrightarrow{W}(i,j,r,\ell)$ and $\overleftarrow{W}(i,j,r,\ell)$ and thus find a w-wavelength virtual topology which is optimal with respect weighted-average hop distance: First, from Theorem 6, $\overrightarrow{W}(i,i,r,\ell)$ and $\overleftarrow{W}(i,i,r,\ell)$ are set to 0 for all i, $0 \leq i \leq n$ and $0 \leq r$, $l \leq w$. Similarly, $\overrightarrow{W}(i,j,0,\ell)$ and $\overleftarrow{W}(i,j,r,0)$ are set to ∞ for all values of i, j, ℓ, and r. Then, for each value of x ranging from 1 to n, for each i, j pair where $0 \leq i, j \leq n$ and $|i-j| = x$, we compute $\overrightarrow{W}(i,j,r,\ell)$ and $\overleftarrow{W}(i,j,r,\ell)$ for all $0 \leq r, \ell \leq w$ as indicated in Theorem 6. Note that this ordering of the computation ensures that $\overrightarrow{W}(i,j,r,\ell)$ and $\overleftarrow{W}(i,j,r,\ell)$ are computed using only previously computed terms.

In Theorem 6, each term $\overrightarrow{W}(i,j,r,\ell)$ and $\overleftarrow{W}(i,j,r,\ell)$ examines $O(n^2)$ combinations of indices s and t. For each s, t pair, the first three terms in the equations $\overrightarrow{W}(i,j,r,\ell)$ and $\overleftarrow{W}(i,j,r,\ell)$ in Theorem 6 are found in constant time by dynamic programming lookup. The fourth terms can also be computed in constant time by performing the following preprocessing step prior to executing the dynamic program: For each i, $1 \leq i \leq n$, compute the partial sum $S(i) = \sum_{j=1}^{i} c(v_j)$ and define $S(0) = 0$. Then, within the dynamic program we can compute $\sum_{k=t}^{j} c(v_k)$ as $S(j) - S(t-1)$, which requires constant time. Using an analogous computation, we can find $\sum_{k=j}^{t} c(v_k)$ in constant time. The preprocessing step takes time $O(n)$. Finally, the number of entries computed by the dynamic program is $O(n^2w^2)$ and each entry takes time $O(n^2)$, resulting in a total running time of $O(n^4w^2)$.

The case of minimizing the weighted-average in bidirectional rings is an immediate extension of this result. By Lemma 2, there exists an optimal solution which is crossing-free. Thus, an optimal solution in the ring partitions the destinations into a contiguous block counter-clockwise with respect to the source and a contiguous block clockwise with respect to the source. Each of these blocks induces a path problem which can be solved using the dynamic programming algorithm above. Consequently, we can examine $O(n)$ partitions of the destinations into two contiguous blocks and solve each of these problems in $O(n^4w^2)$ time, resulting in a $O(n^5w^2)$ time algorithm. The running time can be improved to $O(n^4w^2)$ by the following simple observation: The dynamic programming algorithm for paths can first be performed from the source to all n destination nodes in the clockwise orientation. Then, the algorithm can be performed from the source to all n destination nodes in the counter-clockwise orientation. The dynamic programming tables now contain, among their entries, the best solutions from the source to the first k clockwise and counter-clockwise nodes, for each k, $0 \leq k \leq n$. Therefore, we can now look at the $O(n)$ different partitions, and for each one we can determine the weighted-average hop distance in constant time by examining the appropriate dynamic programming table entries. Thus, the running time is $O(n^4w^2)$.

V. Conclusion

In this paper we have examined the problem of finding virtual topologies which minimize the maximum or average hop distance from a source node to a set of destination nodes. Such a virtual topology can be used for either standard multicast communication, in which the source node sends the same data to all destination nodes, or personalized multicast communication, in which the source node may send different data to different destination nodes.

Although the problem of finding optimal virtual topologies is in general NP-complete, we have shown that a simple $O(n)$ time construction simultaneously minimizes the maximum and average hop distance in paths and rings. Moreover, for the more general case of minimizing the weighted average (or maximum) hop distance in paths and rings, a polynomial-time dynamic programming algorithm can be employed.

Finally, we note that some optical switches have the capability of splitting a lightpath to multiple output ports. Using this functionality, a single lightpath may be "tapped" by multiple nodes. The number of nodes that can tap a lightpath is bounded, however, since the lightpath experiences some loss of power each time it is tapped. The results presented in this paper can be easily generalized for the case that a lightpath can be tapped at any given number of nodes [6]. The generalization of this problem in which each physical link may have a different number of available wavelengths is evidently computationally more complex and some results on this problem have been obtained in [9].

Acknowledgment

The authors wish to thank the anonymous reviewers for making several suggestions which significantly improved this paper.

References

[1] M. Ali and J. Deogun, "Power-efficient design of multicast wavelength-routed networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 1852–1862, Oct. 2000.

[2] D. Banerjee and B. Mukherjee, "Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study," *IEEE/ACM Trans. Networking*, vol. 8, pp. 598–607, Oct. 2000.

[3] R. Dutta and G. N. Rouskas, "On optical traffic grooming in WDM rings," in *Proc. ACM SIGMETRICS/PERFORMANCE*, June 2001.

[4] O. Gerstel, I. Cidon, and S. Zaks, "The layout of virtual paths in ATM networks," *IEEE/ACM Trans. Networking*, vol. 4, pp. 873–884, Dec. 1996.

[5] O. Gerstel, A. Wool, and S. Zaks, "Optimal layouts on a chain ATM network," *Discrete Appl. Math.*, vol. 83, no. 1–3, pp. 157–178, 1998.

[6] J. R. Hartline, R. Libeskind-Hadas, K. Dresner, E. Drucker, and K. Ray, "Optimal virtual topologies for one-to-many communication in WDM paths and rings: Algorithms and intractability results," Dept. Comput. Sci., Harvey Mudd College, Claremont, CA, Tech. Rep. HMC-CS-02-01, Mar. 2002.

[7] W. Hu and Q. Zeng, "Multicasting optical cross connects employing splitter-and-delivery switch," *IEEE Photon. Technol. Lett.*, vol. 10, pp. 970–972, July 1998.

[8] W. Liang and H. Shen, "Multicasting and broadcasting in large WDM networks," in *Proc. 12th Int. Parallel Processing Symp. (IPPS/SPDP)*, 1998, pp. 516–523.

[9] R. Libeskind-Hadas and R. Melhem, "Multicast routing and wavelength assignment in multi-hop optical networks," *IEEE/ACM Trans. Networking*, vol. 10, pp. 621–629, Oct. 2002.

[10] L. H. Sahasrabudde and B. Mukherjee, "Light-trees: optical multicasting for improved performance in wavelength-routed networks," *IEEE Commun. Mag.*, vol. 37, pp. 67–73, Feb. 1999.

[11] H. Shen, F. Chin, and Y. Pan, "Efficient fault-tolerant routing in multihop optical WDM networks," *IEEE Trans. Parallel Distribut. Syst.*, vol. 10, pp. 1012–1025, Oct. 1999.

[12] D. Thaker and G. Rouskas, "Multi-destination communication in broadcast WDM networks: a survey," *Opt. Networks*, vol. 3, no. 1, pp. 34–44, Jan./Feb. 2002.

[13] J. Yates, M. Rumsewicz, and J. Lacey, "Wavelength converters in dynamically reconfigurable WDM networks," *IEEE Commun. Surveys*, Second Quarter 1999.

[14] X. Yuan, R. Melhem, and R. Gupta, "Performance of multi-hop communications using logical topologies on optical torus networks," *J. Parallel Distribut. Comput.*, vol. 61, no. 6, pp. 748–766, June 2001.

[15] X. Zhang, J. Wei, and C. Qiao, "Constrained multicast routing in WDM networks with sparse light splitting," in *Proc. INFOCOM*, 2000, pp. 1781–1790.

Jeff R. K. Hartline received the B.S. degree in computer science and mathematics from Harvey Mudd College, Claremont, CA, in 2001. He is currently working toward the Ph.D. degree in computer science at Cornell University, Ithaca, NY.

His research interests are algorithms, data structures, and information networks.

Ran Libeskind-Hadas (S'91–M'93) received the B.A. degree in applied mathematics from Harvard University, Cambridge, MA, in 1987 and the M.S. and Ph.D. degrees in computer science from the University of Illinois at Urbana-Champaign in 1989 and 1993, respectively.

Since 1993, he has been with the Department of Computer Science, Harvey Mudd College, Claremont, CA, where he is currently a Professor. His research interests are in multicast communication and optical networking.

Dr. Libeskind-Hadas is a Member of the Editorial Board of the IEEE TRANSACTIONS ON COMPUTERS.

Kurt M. Dresner received the B.S. degree in computer science and mathematics from Harvey Mudd College, Claremont, CA, in 2002. He is currently working toward the Ph.D. degree in computer science at the University of Texas, Austin.

His current research interests are in algorithms, artificial intelligence, and complexity.

Ethan W. Drucker received the B.S. degree in computer science from Harvey Mudd College, Claremont, CA, in 2002. He is currently working toward the Ph.D. degree in computer science at the University of California, Los Angeles.

His research interests are in computer graphics, algorithms, and numerical analyisis.

Katrina J. Ray received the B.S. degree in computer science from Harvey Mudd College, Claremont, CA, in 2002. She is currently working toward the Ph.D. degree in computer science at the New Mexico State University, Las Cruces.

Her current research interests are in graph theory and graph algorithms.