Heuristic for Setting up a Stack of WDM Rings with Wavelength Reuse

A. S. T. Lee, D. K. Hunter, *Member, IEEE, Member, OSA*, D. G. Smith, and D. Marcenac

***Abstract*—A heuristic methodology is proposed for the setting up of a stack of wavelength-division-multiplexing (WDM) rings with wavelength reuse when the design traffic exceeds the capacity of a single ring and no wavelength conversion is employed. A *ring stack* consists of an overlay of rings routed over the same physical route, and it can be setup and dimensioned in a myriad of ways. The design traffic comprises of a set of bidirectional *lightpaths* or wavelength connections. There exists a tradeoff between the number of nodes and the number of rings required to carry this traffic, and it is demonstrated that both cannot be minimized simultaneously. For certain traffic patterns, we identify stacks requiring the minimum number of nodes or WADM's, which is desirable from a cost point of view, and stacks requiring the minimum number of rings. An algorithm is presented that manipulates the tradeoff phenomenon to produce a spectrum of designs with deterministic composition. We finally conclude by identifying factors that may influence the choice of design.**

***Index Terms*—Network reliability, synchronous digital hierarchy, topology, wavelength-division multiplexing (WDM).**

I. Introduction

WAVELENGTH-DIVISION-MULTIPLEXING (WDM) optical rings represent the next stage of network deployment for operators needing to meet continuing demand growth. The ring architecture offers inherent and rapid resilience and simpler control and routing than mesh networks. Several designs of WDM optical rings have been classified and are analogs of those based on SONET/SDH-based transmission [1].

Since WDM enhances the transport capability with demonstrations of aggregate throughput of several hundreds of gigabits per second, the total network cost will become increasingly dominated by nodal equipment costs [2]. A scenario that arises is that most of the infrastructure costs will migrate toward the network edge, where most of the cost is concentrated at intelligent high-speed transport network nodes. This will force future designs of broadband networks to gravitate toward those economical solutions where nodal costs (cost of optical add–drop multiplexers (OADM's), switches and electronic multiplexing equipment) are minimized [2]–[4]. This work examines the problem of setting up multiple rings when the design traffic to be carried exceeds the capacity of a single ring. We assume that the line rate and number of wavelengths is fixed. We define a *ring stack* as a collection of rings routed over the same physical route. A heuristic methodology is proposed for setting up a stack of WDM rings with wavelength reuse when given a known set of bidirectional *lightpaths* or wavelength connections, and the capacity of each ring, in terms of the number of supported wavelengths. Ring architectures featuring wavelength reuse include the two-fiber and four-fiber bidirectional WDM rings with shared protection, where a half of the ring capacity is reserved for protection against link and node failures. The stack setup algorithm is also applicable to rings with no protection. The design objective is to minimize the aggregate node count as this is related to the cost of the stack. Several studies [5]–[9] have examined the near term architecture: SONET/SDH rings operating as higher level networks over WDM rings whereby each lightpath in the network is terminated by a SONET/SDH add–drop multiplexer (ADM). The cost of such networks predominantly comprises of ADM costs and significant reductions in the quantity of terminal equipment may be achieved by employing the *optical bypass* facility of wavelength add–drop multiplexers (WADM's). SONET/SDH traffic may be groomed onto individual lightpaths in such a way that the number of ADM's is minimized. However, minimizing the number of ADM's does not necessarily coincide with the minimization of the number of wavelengths [5], [6], [9] in the network. Modiano and Chiu [5], [6] have demonstrated that the general traffic-grooming problem is NP-complete [5] and they have provided algorithms for unidirectional rings. Simmons, Goldstein, and Saleh [7], [8] and Gerstel, Lin, and Sasaki [9] have examined traffic-grooming schemes for bidirectional rings.

This paper is organized as follows. In Section II, we introduce notation and terminology and define the optimization problem to be solved. In Section III we describe a novel heuristic methodology for setting up a ring stack with the main objective of minimizing the aggregate node count. Section IV assesses the results, and conclusions are presented in Section V.

II. Notation and Problem Definition

We distinguish between two types of structure: a stack consisting of equally sized rings is known as a *uniform ring stack* where all nodes physically coincide, and hence all rings are of the same size. Otherwise the structure is referred to as a *variable ring stack*. In the latter each constituent ring in the stack may adopt a different size. Here, a "node" corresponds to an OADM on a particular fiber. It is possible for several nodes to exist at one location, each on separate fibers. As an example, consider rings with up to eight nodes and a maximum of two wavelengths with the following set of demands in Table I to be satisfied. For the uniform ring stack in Fig. 1(a), one lightpath

Manuscript received September 14, 1999; revised January 17, 2000.

A. S. T. Lee, D. K. Hunter, and D. G. Smith are with the Department of Electronic and Engineering, University of Strathclyde, Glasgow G1 1XW, U.K. (e-mail: d.g.smith@eee.strath.ac.uk).

D. Marcenac is with BT Labs, Martlesham Heath, Ipswich IP5 7RE, U.K.

Publisher Item Identifier S 0733-8724(00)03046-2.

0733–8724/00$10.00 © 2000 IEEE

TABLE I
SET OF DEMANDS FOR EIGHT-NODE STACKED RINGS EXAMPLE SHOWN IN FIG. 1

Demand (node pair)	Number of lightpaths
(1, 3)	2
(1, 7)	2
(2, 4)	2
(2, 8)	2
(3, 5)	2
(4, 6)	2
(5, 7)	2
(6, 8)	2

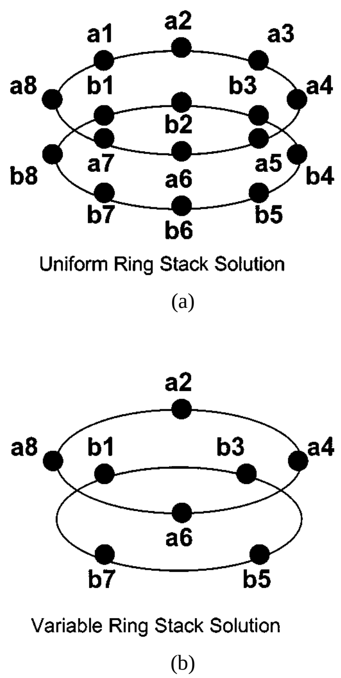

Fig. 1. An example demonstrating reduction in node count (OADM's) by stacking rings.

from each demand is allocated to each eight-node ring, giving a total of 16 OADM's. Alternatively, for the variable ring stack in Fig. 1(b), demands (2, 4), (2, 8), (4, 6), and (6, 8) are routed on ring 1 (nodes a2, a4, a6, and a8) and demands (1, 3), (1, 7), (3, 5), and (5, 7) are routed on ring 2 (nodes b1, b3, b5, and b7), giving a total of 8 OADM's. This example demonstrates that the latter configuration can achieve a lower aggregate node count since the variable ring stack can conform more closely to the traffic pattern.

We assume that each fiber link in the ring supports the transmission of up to λ wavelengths and no wavelength conversion is assumed within the ring. A *ring stack*, S, is the set of rings, $S = \{R_1, R_2, \cdots, R_i, \cdots, R_r\}$ routed over the same physical route, where r is the number of rings in the stack. Each ring R_i is the set of nodes, $R_i = \{a_1, a_2, \cdots, a_{n_i}\}$, $n_i = |R_i|$, $n_{\min} \le n_i \le n_{\max}$, where n_i is the number of nodes in the ring, and $n_{\min}$ and $n_{\max}$ denote the minimum and maximum ring sizes respectively. The total number of nodes in the stack n_S is given by $n_S = |R_1| + |R_2| + \cdots + |R_i| + \cdots + |R_r|$. The general description of an optimization problem can thus be stated as follows.

1) Variable Ring Stack Optimization Problem (VROP): Given a set of lightpaths T between n nodes, a link capacity of λ wavelengths, and the restrictions on ring size, $n_{\min}$ and $n_{\max}$, we wish to find a set of rings S, such that the total number of nodes n_S, in the ring set S is minimized.

Due to the use of a nonuniform ring stack, wavelength reuse may become more efficient for larger capacity networks, hence n_S may not be proportional to T. The predetermined setup traffic consists of c bidirectional lightpaths with each pairwise demand (i, j) comprising k_{ij} lightpaths terminating on any two nodes, i and j, in the node set $I = \{1, 2, \cdots, n\}$, for $0 \le k_{ij} \le k_{\max}$, where $k_{\max}$ is a predetermined ceiling value. Each lightpath, L_k, is represented by its node pair (i, j), for $i, j \in I$, $i < j$. We also define the demand sequence, $\boldsymbol{D}$,

$$\boldsymbol{D} = \{L_1, L_2, \cdots, L_k, L_{k+1}, \cdots, L_c\} \tag{1}$$

as a permutation of the sequence of lightpaths containing all members of the setup traffic set $T = \{k_{ij}$ entries of (i, j) for $i, j \in I, i < j\}$.

III. HEURISTIC METHODOLOGY

We introduce a heuristic methodology for dimensioning a stack of rings given a set of predetermined lightpaths and the number of wavelengths supported on each ring. A routing and wavelength assignment (RWA) heuristic, which we refer to as the cyclic ring RWA algorithm, is employed to produce a near-optimal arrangement of lightpaths requiring the minimum number of wavelengths on the ring. For a full mesh of lightpaths (one connection for each unique node pair), optimal RWA algorithms have been developed by several authors [10]–[17]. These algorithms all specify an RWA system in which the set of lightpaths is partitioned into subsets. Each constituent subset of lightpaths is then assigned to a unique wavelength. Different optimal RWA configurations have been discovered with each type of configuration identified by its method of partitioning the full mesh of lightpaths. For arbitrary traffic patterns (no more than one lightpath per unique node pair, i.e., $k_{ij} = 0$ or $k_{ij} = 1$), we present a simple heuristic for seeking near-optimal RWA configurations.

A. Setting Up the Uniform Ring Stack (UNn)

The following presents a scheme for setting-up the uniform ring stack. Partition an m-wavelength, n-node base ring into constituent $\lceil m/\lambda \rceil$ λ-wavelength, n-node rings, where m is the minimum number of wavelengths for an optimal routing and wavelength assignment (RWA) to support T on an n-node ring. This stack has the maximum average ring size.

B. Setting Up the Variable Ring Stack (VR2)

This is the case where each constituent ring has the minimum size of two nodes. This stack is setup as follows. Partition T into n_U subsets, with each subset U_i, $1 \le i \le n_U$ containing identical lightpaths having identical endpoints in all cases. The number of rings in the stack is given by $r_R = \sum_i \lceil |U_i|/2\lambda \rceil$ with $n_S = 2r_R$. This stack has the minimum average ring size and requires a minimal structure of at least one ring for every unique node pair in T.

C. Heuristic for Setting Up Variable Ring Stacks

The VROP can be decomposed into two optimization subproblems: 1) identifying a stack of rings with a minimum aggregate node count and 2) for each ring, finding an optimal RWA configuration. The decomposition approach to the problem cannot guarantee a solution that is globally optimal. The variable ring stack is evolved from an initial two-node ring ($n_{\min} = 2$) or "seed" and augmented by the serial addition of lightpaths to the stack. The order in which lightpaths are allocated is specified by the demand sequence $\boldsymbol{D}$. A formal description of the algorithm is given as follows.

Step 1 (Create the "Seed"): A two-node ring R_1 is created to accommodate the first lightpath, L_1 in $\boldsymbol{D}$, with the ring nodes corresponding to the terminations of the lightpath connection. The allocation of L_1 to R_1 is fixed and an arbitrary route and wavelength is chosen since the configuration may change in successive steps.

Step 2 (Accommodate Successive Lightpaths): The next lightpath in the sequence, L_k, is tentatively allocated to a ring R_i, $1 \leq i \leq r$ in an evolving stack of currently r rings. The RWA configuration in ring R_i is recomputed. If the new RWA configuration can be supported on this ring, the lightpath is permanently assigned to this ring, otherwise the lightpath is allocated to another ring in the stack. If no ring can be found, a new two-node ring R_{r+1} is created to accommodate this lightpath. Note that if the evolving stack comprises more than one ring ($r > 1$), a choice must be made in selecting a ring for the tentative allocation of a lightpath. Two ring-selection rules considered in this paper are the following.

1) First Fill Scheme (FF): Ring candidates $\{R_i, 1 \leq i \leq r, r$ rings in stack $\}$ are considered in chronological order, that is, the algorithm will attempt to allocate the current lightpath, L_k, to the first ring created, i.e., R_1. Failing that, it will try the next "oldest" ring in sequence, R_2, and so on, until a ring is found, or a new ring R_{r+1} must be created to accommodate the lightpath.

2) Minimum Fill Scheme (MF): Ring candidates $\{R_i, 1 \leq i \leq r, r$ rings in stack$\}$ are considered in descending order of the number of node matches, $E_i = |L_k \cap R_i|$, between ring R_i and current lightpath L_k. The first candidate chosen is the ring with the greatest matching of nodes, i.e., the ring with the highest E_i value, with ties broken chronologically. If no ring can be found, a new ring R_{r+1} is created to accommodate the lightpath. This scheme favors filling those rings that exhibit minimal growth upon addition of a lightpath.

A flow chart representation of the scheme is shown in Fig. 2.

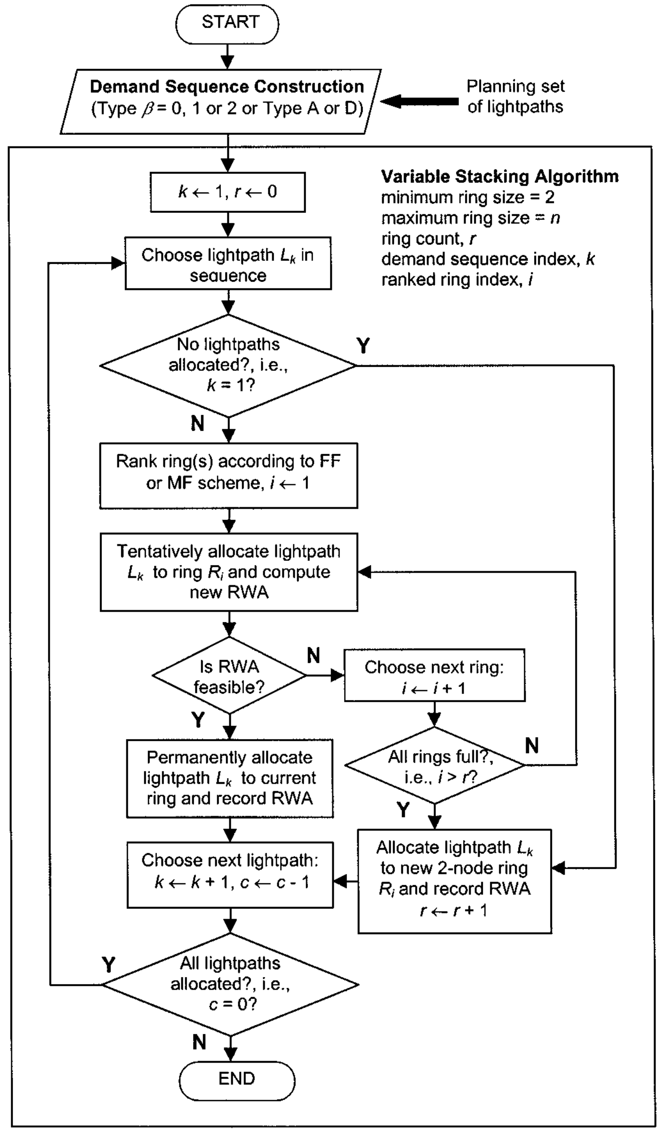

Fig. 2. Flow diagram representation of the variable ring stacking algorithm.

D. Demand Sequence Generation

For the algorithm described in Section III-C, a ring is extended if $|L_k \cap R_i| < 2$ and remains constant in size if $|L_k \cap R_i| = 2$. As an example, consider the case where the first two lightpaths, $L_1 \equiv (s_1, t_1)$ and $L_2 \equiv (s_2, t_2)$ are allocated to the first two-node ring R_1. On allocation of lightpath L_2 to ring R_1, R_1 only doubles in size if node indexes s_1, s_2, t_1, and t_2 are unique. If $s_1 = s_2$ and $t_1 = t_2$, the ring is not extended as no new traffic sources are added to the ring. Thus, we can choose specific permutations of $\boldsymbol{D}$ to "control" the growth rate of rings and the rate at which they are filled. Since an exhaustive assessment of all possible $c!$ permutations is not feasible for a sequence of c lightpaths, we examine only a small subset of permutations for generating structures spanning the two design extremes represented by the UN8 and VR2-type designs.

1) Type H β Sequence: β is the number of times that each pair of adjacent lightpaths, (s_k, t_k) and (s_{k+1}, t_{k+1}) in $\boldsymbol{D}$ have matching node indexes over all k

$$\beta = \begin{cases} 0, & s_k \neq s_{k+1} \text{ and } t_k \neq t_{k+1} \\ 1, & s_k \neq s_{k+1} \text{ and } t_k = t_{k+1} \\ & \text{or } s_k = s_{k+1} \text{ and } t_k \neq t_{k+1} \\ 2, & s_k = s_{k+1} \text{ and } t_k = t_{k+1} \end{cases}. \tag{2}$$

An Hβ type sequence is one in which every adjacent pair of lightpaths in the sequence satisfies the constraint indicated by the value for β. Table I summarizes the salient stack evolution characteristics for each value of β. Observe that the H0 type

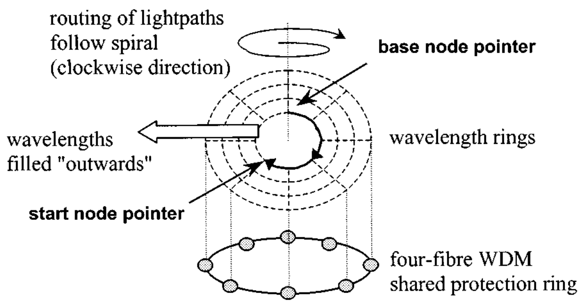

Fig. 3. "Snapshot" of the cyclic ring RWA algorithm.

sequence favors the creation of rings with the maximum size whilst the H2 type sequence will produce a high population of two-node rings in the stack. The H1 type sequence produces a stack that is "structurally between" H0 and H2 generated designs. The construction of $\boldsymbol{D}$ is initiated by randomly choosing a lightpath and successive lightpaths are added to the growing sequence in a greedy fashion. We consider each lightpath in random order and concatenate to $\boldsymbol{D}$ if it satisfies the constraint indicated by the value for β or attempt the next lightpath in sequence if the constraint is violated. Lightpaths added to $\boldsymbol{D}$ are removed from the preparatory sequence of unconcatenated lightpaths. If no appropriate lightpath can be found, i.e., the end of the preparatory sequence is reached, we append the first lightpath to $\boldsymbol{D}$ and restart the construction process for successive lightpaths until all have been incorporated in $\boldsymbol{D}$. Note that the uniformity of the constraint along $\boldsymbol{D}$ may not be possible for a particular randomly generated demand set.

2) Type D and A Sequence: Lightpaths are sorted in order of descending or ascending hop count respectively to produce these sequences, i.e., assuming the shortest path on an n node ring. Note that the sorting process brings together lightpaths with identical node pairs and as a consequence the sequence can be described as a concatenation of H0-type subsequences.

E. Cyclic Ring Routing and Wavelength Assignment

We consider a ring with n nodes with a clockwise numbering of nodes along the ring. The ring RWA algorithm employs a system of two pointers for routing lightpaths along a "spiral" from the innermost wavelength (a "snapshot" of the algorithm is depicted in Fig. 3), assuming an ordinal numbering of wavelengths. The direction of the "spiral" is arbitrary and taken to be in the clockwise direction. The two pointers are referred to as the *base node* n_b and the *start node* n_s and both are updated in the same direction as the algorithm progressively seeks available lightpaths terminating at n_s to fill wavelengths along the spiral. A formal description of the algorithm is given below and valid for both clockwise and counter-clockwise operations of the algorithm.

Step 1 (Initialization): Pointers n_b and n_s are initially assigned to a chosen node on the ring. The current wavelength w is initialized to wavelength 1.

Step 2 (Find Lightpath Terminating at n_s): A lightpath (s, t) terminating at n_s is sought from the set of unallocated lightpaths, and satisfies the following conditions: a) the shortest path must be in the clockwise direction along the ring; i.e., node pair either of the form (n_s, s) or (n_s, t) and b) the path must fit in s contiguous unused wavelength slots at wavelength w, starting at n_s and ending at n_b in the clockwise direction. If n is even, for lightpaths of length $n/2$ there exist two paths of equal length. The route chosen is dictated by the directional operation of the algorithm, i.e., assuming a clockwise operation and a lightpath between nodes n_s and s, the clockwise path (n_s, s) is always chosen.

The number of unused wavelength slots between two consecutive nodes is given by

$$s = \begin{cases} n, & \text{if } n_s = n_b \\ |n_b - n_s|, & \text{otherwise.} \end{cases} \tag{3}$$

If a lightpath is found that satisfies the conditions, it is routed on wavelength w and n_s is updated to the other node that terminates the lightpath. If no lightpath can be found which terminates at n_s or n_b, the pointer is advanced one node position clockwise and the search resumes at this new position. If both pointers are realigned, a cycle is completed at wavelength w and the next wavelength is chosen: $w \leftarrow w + 1$. The number of wavelengths required to support the lightpaths, w', is equal to the final value for w.

Note that only shortest paths are considered and as a consequence an optimal configuration which requires the minimum number of wavelengths is not guaranteed since such a configuration may require a subset of lightpaths to be routed along their longest path. The algorithm is executed n times, i.e., once for each of the n possible starting positions n_b in the ring. The solution with the lowest wavelength requirement, $\min_n w'$, is chosen as the final solution.

IV. Results and Discussion

The average values for the aggregate node count $\overline{n}_S$, ring count $\overline{r}_R$, and the ring size $\overline{n}_{R_i}$ are determined from multiple traffic sets [the number of lightpaths, $k_{ij} \in (0, k_{\max})$ per demand (i, j), $i, j \in I$, $i < j$ is uniform randomly distributed] for each type of stack. In addition each traffic set is scaled up to ten times the volume.

A. Comparing Demand Sequence Types

Combining a demand sequence with the FF or MF variants of the heuristic produces a unique setup algorithm for the variable ring stack. The symbol for each constituent component is concatenated to identify each unique combination. For example, when the H1 type sequence is coupled with the FF scheme, we refer to this as theFF-H1 method. Multiple traffic sets are randomly generated for $k_{\max} = 5$ and $n = 8$ for a comparison of all possible setup algorithms. Graphs of $\overline{n}_S$, $\overline{r}_R$, and $\overline{n}_{R_i}$ are plotted against the traffic scaling factor and depicted in Figs. 4–6, respectively, with curves for the UN8, VR2, and variable ring stacks plotted on the same graph.

From Fig. 4, we observe that on average, the variable ring stack achieves a lower $\overline{n}_S$ than the VR2 structure when the traffic scaling is $2\times$ or below. Compared to the other stacks, the VR2 stack requires a minimum stacking structure of at least a two-node ring per unique node pair. This represents an over-

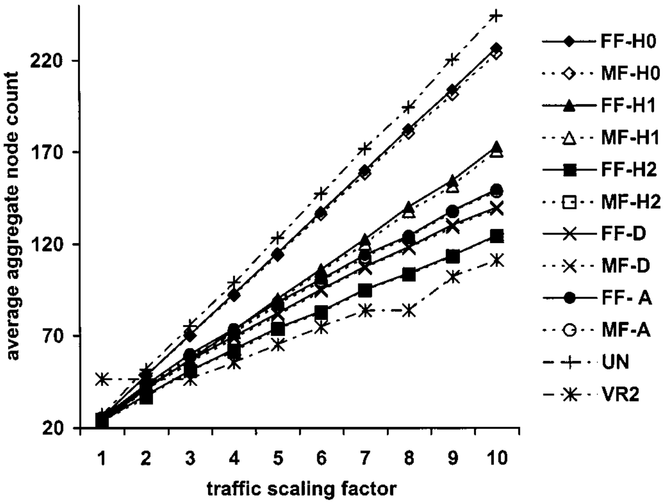

Fig. 4. Plot of the average aggregate node count against the traffic volume.

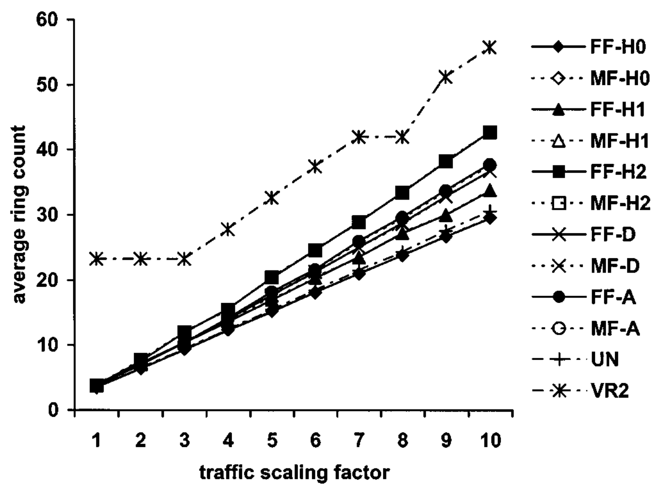

Fig. 5. Plot of the average ring count against the traffic volume.

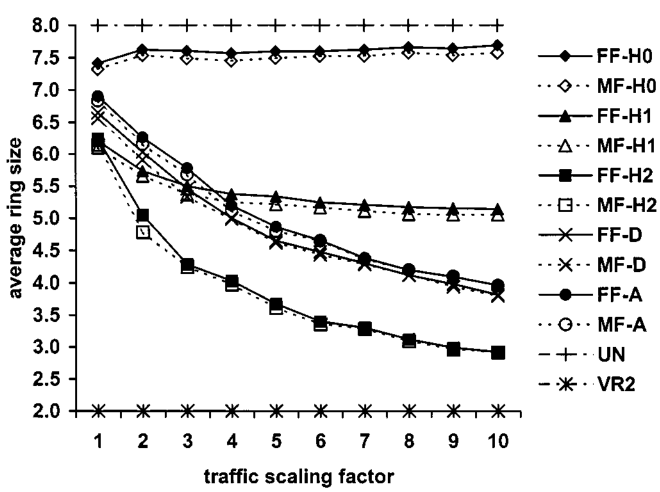

Fig. 6. Plot of the average ring size versus traffic volume.

dimensioning due to a greater number of underutilized wavelengths. However, for higher traffic volumes, the VR2 stack yields the lowest $\overline{n}_S$. The VR2 stack is more efficient as it can adapt more closely to the traffic pattern thus minimizing the

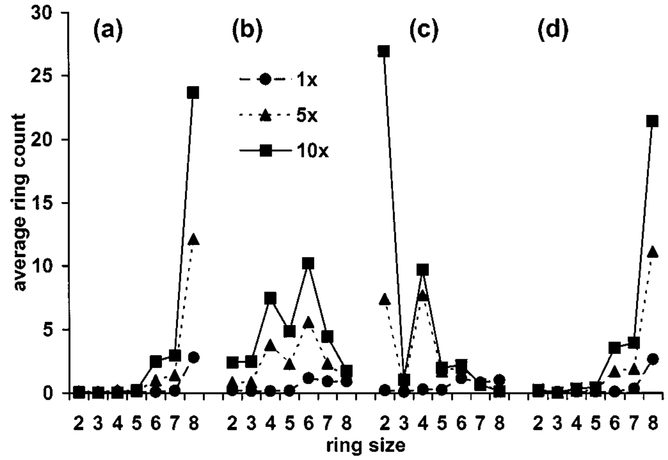

Fig. 7. Distribution of ring sizes comprising the stack. For each method (a) MF-H0, (b) MF-H1, (c) MF-H2, and (d) MF-D and for $1\times$, $5\times$, and $10\times$ scaling of random traffic.

number of unused wavelengths in the stack. The highest $\overline{n}_S$ is exhibited by the UN8 scheme as it provides the coarsest dimensioning. In contrast, for the variable ring stack, the node reductions attained is attributed to the cumulative effect of minimally augmenting the structure of the stack whenever a lightpath is added. For the FF scheme, we minimize the addition of a new ring to the stack at each stage while in the MF scheme we minimize the number of nodes added at each stage. Comparing the FF and MF schemes, the MF scheme on average produces a stack with the lower $\overline{n}_S$ of the two. This is because of the finer resolution of augmentation during stack evolution. However, this gain is marginal for low volumes of setup traffic and only improves when the traffic is sufficiently high.

Comparing Figs. 4 and 5, it is evident that a tradeoff exists between $\overline{n}_S$ and $\overline{r}_R$. The VR2 structure has the highest $\overline{r}_R$ and the lowest $\overline{n}_S$, requiring more than twice the number of rings than the FF-H0 method, which produces a stack with the lowest $\overline{r}_R$. Although the H0-FF scheme on average yields the lowest $\overline{r}_R$ we would expect it to exhibit the highest $\overline{n}_S$.

On the other hand, the MF-H2 method produces a stack with the lowest aggregate node count. It is apparent that the goal of minimizing the aggregate node count does not coincide with the minimization of the ring count. As $\overline{n}_{R_i}$ and $\overline{n}_S$ decrease, $\overline{r}_R$ increases since ring capacity is proportional to ring size. More rings are required for stacks with a low average ring size, as more of them are required to satisfy the same volume of demands. Also note that only the FF-H0 and MF-H0 methods, on average, produce a stack with a lower $\overline{r}_R$ than that of the uniform case.

In Fig. 7(a)–(d), a set of ring population curves (for each size of ring from $n = 2$ to $n = 8$) is plotted for the FF scheme and for H0, H1, H2, and D types of demand sequences, respectively. The set of curves for type A is very similar to the type D sequence and is not shown. These curves depict the rise and stabilization of populations of rings of certain sizes with increasing traffic volume. Each type of sequence favors the dominance of a particular subset of ring sizes and is intimately related to the property of the demand sequence. For sequences of type H1, H2, D, and A, $\overline{n}_{R_i}$ decreases with increasing traffic. Inspecting

TABLE II
COMPARING STACK EVOLUTION CHARACTERISTICS FOR TYPE H β DEMAND SEQUENCE

n	Ring growth rate	Ring fill rate	Most common ring size
0	*Fast*	*Slow* –large rings take longer to fill	n_{max} –rings have a better chance to reach their maximum size
1	*Moderate*	*Moderate*	A range of sizes in $\{n_{min}, n_{max}\}$
2	*Zero* to *Slow*	*Fast* –small rings are filled quickly	n_{min} –most newly created rings are filled continuously

TABLE III
RESULTS FOR FF-H β METHODS AT $5\times$ AND $10\times$ SCALING OF RANDOM TRAFFIC

	H0		H1		H2	
	5x	10x	5x	10x	5x	10x
Avg. aggregate node count	115.93	226.53	90.07	175.16	74.67	123.75
Avg. ring count	15.15	29.51	16.9	33.79	20.34	42.66
Avg. ring size	7.65	7.68	5.33	5.18	3.67	2.9
Most common ring size	8	8	6	6	2	2

Fig. 7(b)–(d) reveals that this phenomenon is caused by a population shift of rings from sizes equal or close to $n_{\max}$, to rings at sizes below this. For the type H0 sequence, $\overline{n}_{R_i}$ is constant as the dominance of rings with size $n_{\max}$ is maintained for all traffic volumes.

Selected results are presented in Table III and comparing this to Table II, we demonstrate that designs are produced with features that have been predetermined. The following section demonstrates that evolution and composition of the stack is deterministic and both are intimately linked.

B. *Deterministic Stack Evolution and Composition*

Fig. 8(a)–(c) depicts the average stack composition curves for the MF-H0, MF-H1, and MF-H2 methods respectively and curves are plotted for each value of the traffic scaling. The average ring size of each constituent ring in the stack is determined and rings are ordinarily numbered according to creation order. In each case the curves broaden with increasing traffic, as more rings are required in the stack to carry the traffic. A flattening of the curves is also observed which is attributed to the increasing dominance of a particular subset of ring sizes appearing in the ring population [Fig. 8(a)–(c)]. We distinguish the three deterministic profiles in the following.

1) H0 Type Profile and Discussion on Effects Due to a Nonuniform $\boldsymbol{D}$: This refers to the average stack composition curve of the type H0 sequence shown in Fig. 8(a). The corresponding distribution of ring sizes is shown in Fig. 7(a).

We explain at this juncture how this nonuniformity can influence the evolution of the stack in subtle ways. An imbalance that frequently occurs in some permutations of $\boldsymbol{D}$ is that the construction of the head-end of $\boldsymbol{D}$ exhausts certain node pairs for a particular value of β, e.g., $\beta = 0$, and as a consequence, other properties become more prominent toward the tail end, e.g., $\beta = 1, 2$. For example, consider the H0 type sequence (refer to Table II): the head end of the sequence will exhibit a fast ring growth rate (low ring fill rate) but toward the tail end a slower ring growth rate occurs (also accompanied by a quickening of ring fill rate). This affects the composition of the stack: Large rings appear in the first half of the stack construction phase and

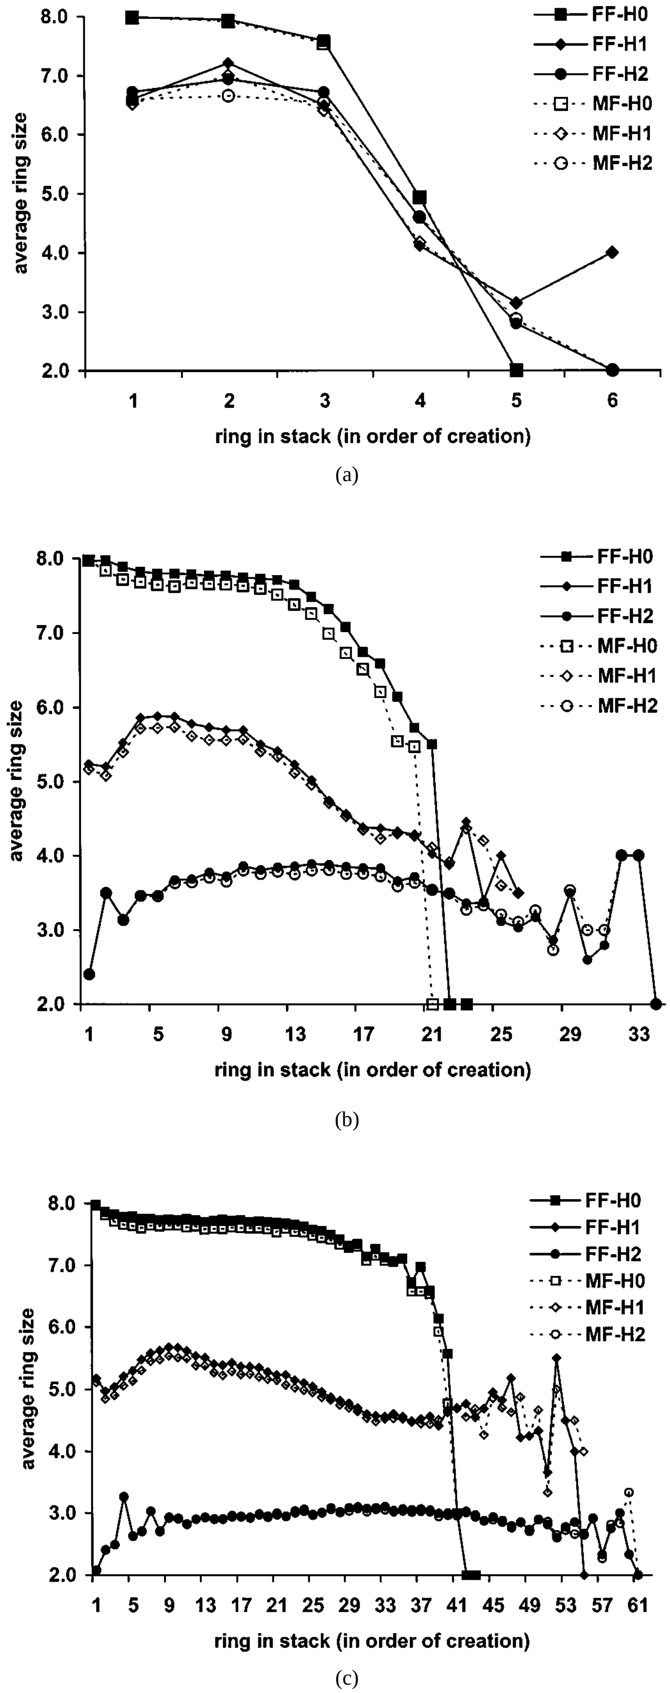

Fig. 8. (a) Average stack composition ($1\times$ random traffic) for methods X-H0, X-H1, X-H2 for X $= \{\mathrm{FF}, \mathrm{MF}\}$, (b) average stack composition ($5\times$ random traffic) for methods X-H0, X-H1, X-H2 for X $= \{\mathrm{FF}, \mathrm{MF}\}$, and (c) average stack composition ($10\times$ random traffic) for methods X-H0, X-H1, X-H2 for X $= \{\mathrm{FF}, \mathrm{MF}\}$.

small rings appear in the latter half. However, the broadening of all profiles with traffic volume indicates that these influences are swamped as the adherence to a value of β improves.

2) H2 Type Profile: This refers to the common profile shared by type H2, D, and A sequences [Fig. 8(c)—The profiles for type D and A sequences are not shown but bear many similarities with that of the type H2 sequence.] The size of the "first ring" decreases rapidly and there is an observed flattening of the profile with increasing traffic volume. This is attributed to the increasing lengths of zero ring growth (fast ring fill rate) periods which has the effect of rapidly filling newly created rings and swelling the population of two-node rings [Fig. 7(c)]. As a consequence, the creation of rings of other sizes is moderated or suppressed. For A and D types of demand sequence, the sorting produces an effect that gives rise to a secondary population of four-node rings. The peaking occurring later in the profile is attributed to the effects of the nonuniformity described earlier.

3) H1 Type Profile: This refers to the average stack composition curve for the type H1 sequence [Fig. 8(b)]. This "hybrid profile" reflects characteristics pertaining to the former two. There is a flattening of the profile toward an average ring size of $n_{\max}/2$: the population of rings with sizes falling in a range bordering this value becomes prominent with increasing traffic volume [Fig. 7(b)]. The $\beta = 1$ property favors the creation of rings with sizes around $n_{\max}/2$ and rings with sizes equal or close to $n_{\min}$ and $n_{\max}$ tend to be suppressed by the effect of moderate ring growth and ring fill rates.

C. Changing the Ring Size and Ring Capacity

$\overline{n}_S$ and $\overline{r}_R$ are computed for n = 4, 6, 8, and 12 (traffic sources) and k = 4, 8, 12, 16 (wavelengths). UN8, VR2, FF-D, MF-D stack setups are applied to multiple uniform random traffic sets for $k_{\max} = 10$. Plots of $\overline{n}_S$ and $\overline{r}_R$ versus k are shown in Figs. 9 and 10, respectively with a curve plotted for each value of n. For all values of n, $\overline{n}_S$ and $\overline{r}_R$ decreases with increasing k as a greater volume of lightpaths can be packed into rings. The VR2 stack achieves the lowest $\overline{n}_S$for low k, but this advantage disappears with increasing k as the minimal structure becomes increasingly inefficient as a high proportion of the stack (wavelengths) is unused. This is also true for the UN8 stack for increasing n and decreasing k, as the dimensioning becomes coarser. Comparing stack structures we observe a tradeoff between $\overline{n}_S$ and $\overline{r}_R$. The MF scheme, on average, performs marginally better than the FF scheme in terms of $\overline{n}_S$ but this is at the expense of a greater $\overline{r}_R$.

D. Summary

To assess how the tradeoff between $\overline{n}_S$ and $\overline{r}_R$ can aid in the selection of a stack design, a plot of $\overline{n}_S$ against $\overline{r}_R$ is produced for the FF scheme and shown in Fig. 11. Observe that data at equal volumes of traffic approximate a curve, with three of these being highlighted (data is delineated for $1\times$, $5\times$, and $10\times$ scaling of traffic). Each curve describes a tradeoff region between $\overline{n}_S$ and $\overline{r}_R$. Two straight lines connecting VR2 and UN8 data are also highlighted. These demarcate the region of all feasible stacks from invalid designs: a stack with an average ring size less than 2 and greater than $n_{\max}$ being impossible.

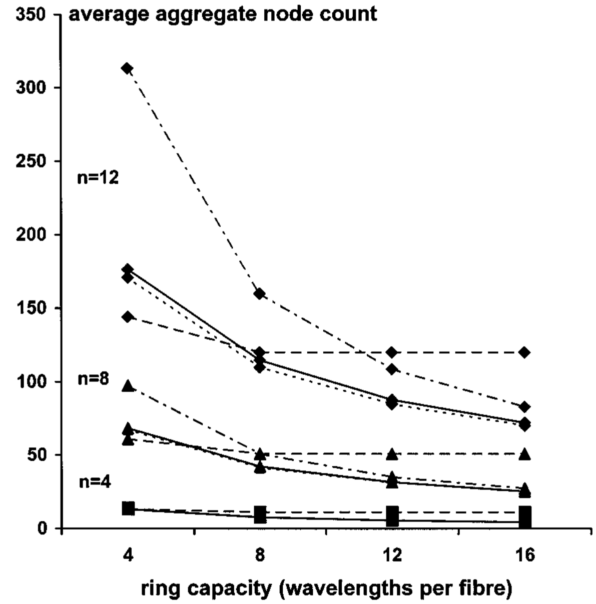

Fig. 9. Plot of $\overline{n}_s$ against $k(k_{\max} = 10)$. For each $n = 4, 8, 12$: $k = 4, 8, 12, 16$ wavelengths and for each method FF-D, MF-D, VR2, and UN.—— FF-D, · · · MF-D, - - - VR2, and -·-·- UN.

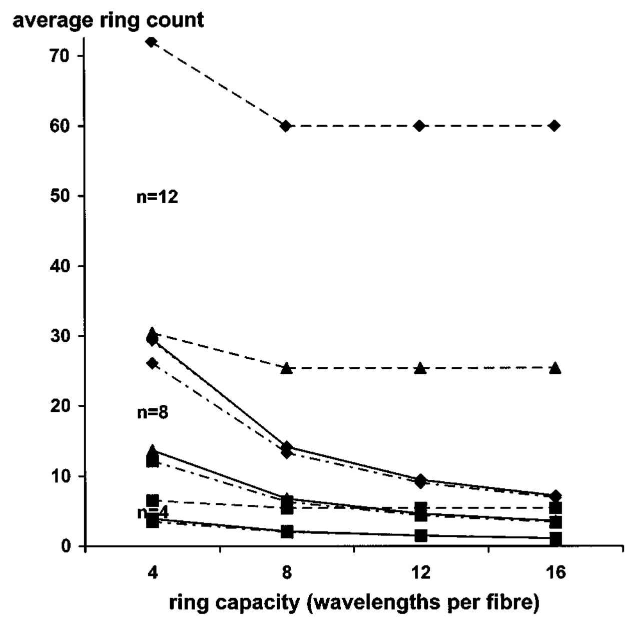

Fig. 10. Plot of $\overline{r}_S$ against $k(k_{\max} = 10)$. For each $k = 4, 8, 12$: $k = 4, 8, 12, 16$ wavelengths and for each method FF-D, MF-D, VR2, and UN.—— FF-D, · · · MF-D, - - - VR2, and -·-·- UN.

The curves of equal traffic may vary for different traffic patterns and we note that for certain traffic patterns the VR2 and UN8 do not represent optimal structures as this requires a variable ring stack design. The curves broaden with increasing traffic volume, as the number of possible designs becomes greater. The implications of the tradeoff between $\overline{n}_S$ and $\overline{r}_R$ will affect the choice of design. The relative proportion of infrastructure costs in setting up the stack: ring placement costs (construction costs, conduits, ducts etc.), link costs (fiber, line terminating equipment, "lighting costs") and node costs (switches, OADM's, concomitant equipment and software) will affect the basic cost of each design. For example, ring placement costs might be artificially

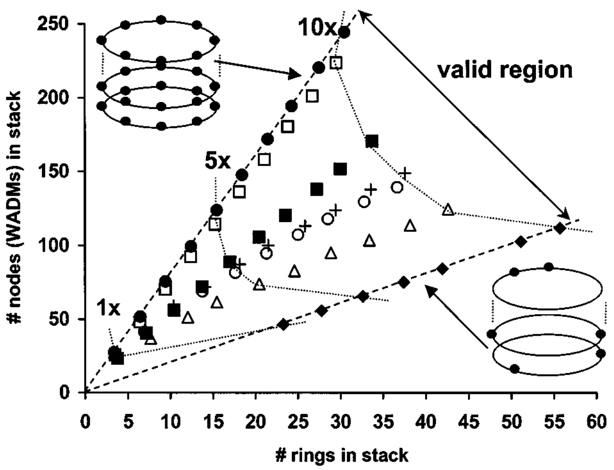

Fig. 11. Plot of number of nodes ($\overline{n}_S$) against number of rings ($\overline{r}_S$) in the stack for each scaling of random traffic (up to $10\times$ traffic). Symbols are □ MF-H0, ■ MF-H1, △ MF-H2, + MF-A, ∘ MF-D, ♦ VR2, and • UN.

high in metropolitan areas where planning is complex and space is at a premium. The deployment of a stack with a high $\overline{r}_R$ may not be desirable in such a scenario. Equipment costs will also vary with manufacturers and implementation, with costs decreasing over time as technology matures in manufacturing, fabrication, and design. However, if the following trends are extrapolated into the future, nodal costs will predominate in the near to short term as more traffic will be aggregated by improvements in WDM and future OTDM-based transmission techniques. We summarize by presenting several factors that may influence the choice of a feasible stack design.

- *Infrastructure Cost*: A stack with a high $\overline{r}_R$ and low $\overline{n}_S$ (VR2, H2-MF, H2-FF) is desirable if node costs are high and ring placement costs are low. However the excessive ring counts may result in a design compromise as we trade for designs with a moderate ring counts at the expense of increasing $\overline{n}_S$ from the minimum structure. A similar scenario exists for stacks with a low $\overline{r}_R$ and high $\overline{n}_S$ (UN8, H0-MF, H0-FF) when node costs are low and ring placement costs are high which is a reasonable assumption in urban areas.
- *Upgrade Complexity*: The composition of the stack may influence the choice of design when the stack has been setup and further traffic is to be carried. Stacks consisting of a large number of small rings are favored as dimensioning these rings is simpler and the rearrangement of lightpaths to pool capacity has a lower impact. If we adhere to minimal augmentations, as this is related to cost, we can upgrade the existing structure via minimal node and ring additions thus the stack can be cost-effectively augmented to meet further growth in demand. Dimensioning a variable ring stack is more complex than the VR2 and UN8 structures, since in both cases an identical ring is simply added when there is insufficient capacity in the stack.

V. Conclusion

A heuristic methodology is presented which confirms the possibility of the cost-effective setup of stacked ring networks. A primary advantage of the setup algorithms is that deterministic structures are produced and the generic framework of the methodology can be adapted to examine a greater number of designs. The tradeoff between the aggregate node count and the number of rings required to implement the stack was quantified for both uniform and variable ring stacks, and this phenomenon is manipulated to produce a spectrum of deterministic structures. These heuristics offer simple real world implementations and their fast execution times suggest their implementation as frequently called routines in network design tools.

References

[1] T.-H. Wu, *Fiber Network Service Survivability*. Norwood, MA: Artech House, 1992.

[2] K. Sato, *Advances in Transport Network Technologies: Photonic Networks, ATM, and SDH*. Norwood, MA: Artech House, 1996.

[3] K. Bala, R. H. Cardwell, D. Feddor, K. Goel, T. Hodgson, H. Kobrinski, M. Mihail, G. Pearson, O. J. Wasem, and M. Wilder, "WDM network economics," in *Proc. NFOEC '96*, Denver, CO, 1996.

[4] S. Johansson, A. Manzalini, M. Giannoccaro, R. Cadeddu, M. Giorgi, R. Clemente, R. Brändström, A. Gladisch, J. Chawki, L. Gillner, P. Öhlén, and E. Berglind, "A cost-effective approach to introduce an optical WDM network in the metropolitan environment," *IEEE J. Select. Areas Commun.*, vol. 16, no. 7, pp. 1109–1122, Sept. 1998.

[5] E. H. Modiano and A. L. Chiu, "Traffic grooming algorithms for minimizing electronic multiplexing costs in unidirectional SONET/WDM ring networks," in *Proc. CISS'98*, Princeton, NJ, Mar. 1998, pp. 653–658.

[6] A. L. Chiu and E. H. Modiano, "Reducing electronic multiplexing costs in unidirectional SONET/WDM ring networks via efficient traffic grooming," in *Proc. IEEE GLOBECOM'98*, Sydney, Australia, Nov. 8–12, 1998.

[7] J. M. Simmons, E. L. Goldstein, and A. A. M. Saleh, "On the value of wavelength-add/drop in WDM rings with uniform traffic," in *Proc. OFC'98*, San Jose, CA, Feb. 22–27, 1998, pp. 361–362.

[8] ——, "Quantifying the benefit of wavelength add–drop in WDM rings with distance-independent and dependent traffic," *J. Lightwave Technol.*, vol. 17, pp. 45–57, Jan. 1999.

[9] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in WDM rings to minimize system cost instead of number of wavelengths," in *Proc. INFOCOM'98*, San Francisco, CA, Mar. 29–Apr. 2 1998.

[10] G. Wilfong, "Minimizing wavelengths in an all-optical ring network," in *Proc. 7th Int. Symp. Algorithms Computation*, 1996, pp. 346–355.

[11] J.-C. Bermond, L. Gargano, S. Perennes, A. A. Rescigno, and U. Vaccaro, "Efficient collective communication in optical networks," in *Proc. 23rd Int. Colloquium on Automata, Languages and Programming—ICALP '96*, Paderborn, Germany, July 1996, pp. 574–585.

[12] C. Qiao and X. Zhang, "Optimal design of WDM ring networks via resource balance," in *IEEE/LEOS Broadband Optical Networks*, Aug. 1996.

[13] C. Qiao, X. Zhang, and L. Zhou, "Scheduling all-to-all lightpaths in WDM rings," in *SPICE Proc. All-Optical Commun. Syst.: Arch., Contr. Network Issues II*, Boston, MA, Nov. 20–21, 1996, pp. 218–229.

[14] G. Ellinas, K. Bala, and G.-K. Chang, "Scalability of a novel wavelength assignment algorithm for WDM shared protection rings," in *Proc. OFC '98*, San Jose, CA, Feb. 22–27, 1998, pp. 363–364.

[15] D. K. Hunter and D. Marcenac, "Optimal mesh routing in four-fiber WDM rings," *Electron. Lett.*, vol. 34, no. 8, pp. 796–797, Apr. 1998.

[16] G. Ellinas, K. Bala, and G.-K. Chang, "A novel wavelength assignment algorithm for 4-fiber WDM self-healing rings," in *Proc. ICC'98*, Atlanta, GA, June 7–11, 1998.

[17] X. Zhang and C. Qiao, "On scheduling all-to-all personalized lightpaths and cost-effective designs in WDM rings," IEEE/ACM Trans. Networking, to be published.

A. S. T. Lee received the B.Eng. degree (Hons) in electronic and electrical engineering from the University of Strathclyde, Glasgow, U.K., in 1997. He is currently working towards the Ph.D. degree from the same university.

His research interests are in the design, analysis, routing, and resilience aspects of all-optical networks.

D. K. Hunter (S'88–M'90) was born in Glasgow, U.K., in 1965. He received the B.Sc. degree in electronic and microprocessor engineering with first class honors from the University of Strathclyde, Glasgow, in 1987. In 1991, he received the Ph.D. degree from the same university for work on optical TDM switching architectures.

Since 1988, he has been with the University of Strathclyde and is now a Senior Research Fellow. Since 1995, he has held an EPSRC Advanced Fellowship. He has visited both BT Labs, U.K., and the University of Pittsburgh, PA, for extended research assignments. His research has involved optical networking, particularly WDM rings and optical packet switching, and is also interested in network survivability, WDM routing and optical network management. He has authored or coauthored over 70 journal and conference publications, and is an active reviewer for many international professional journals and conferences.

Dr. Hunter is an Associate Editor of the IEEE TRANSACTIONS ON COMMUNICATIONS, and is a Guest Editor of the IEE/OSA JOURNAL OF LIGHTWAVE TECHNOLOGY. He is an associate member of the Institution of Electrical Engineers (IEE) and a member of the Optical Society of America (OSA).

D. G. Smith is Professor of Communications Engineering in the Department of Electronic and Electrical Engineering at the University of Strathclyde. His interests are in traffic management, performance evaluation of broad-band networks, and in protection and resilience in networks. Recently, he has been concerned with mechanisms for the efficient transport of MPEG coded video across ATM networks He is closely involved with the work of the Network Management and Resource Centre at Strathclyde, U.K., which has a particular interest in the interaction between network technology and economic, regulatory, and public policy issues.

D. Marcenac was born in Grenoble, France, in 1968. He received the B.A. degree in engineering from St. Catherine's College, Cambridge, U.K., in 1990, and the Ph.D. degree in optoelectronics in 1994 from the same institution.

Since 1994, he has worked with BT Labs, U.K., first, on research in optics involving wavelength conversion and applications of semiconductor optical amplifiers, and more recently in transport network design. He is now part of the global transport network planning and design team for BT Worldwide, designing BT's pan-European and Asia Pacific high-capacity SDH and WDM networks.