Survivable Passive Optical Networks Based on Arrayed Waveguide Grating Architectures

Jane M. Simmons, Senior Member*, IEEE*

***Abstract*— By arranging a small number of AWGs in survivable topologies with specific interconnection patterns, we can construct passive optical networks with a high degree of protection. A variety of designs are presented that differ in the level of survivability, the protection mechanism, and the amount of required end-user equipment.**

***Index Terms*— Arrayed waveguide grating (AWG), passive optical networks (PON), protection, self-healing networks, wavelength grating router (WGR)**

I. INTRODUCTION

PASSIVE optical networks are well suited for regions that are relatively small in geographic extent and number of users. The lack of active components in the network enables rapid deployment, operational simplicity, and reduced maintenance requirements. Arrayed waveguide gratings (AWGs) (also known as wavelength grating routers) are often used as the central routing device in such networks due to their excellent wavelength re-use properties [1][2]. While many of the applications for AWG-based passive networks focus on access, where the end users are communicating with a central office, this paper addresses communication among a small number of users, where any two users potentially need to communicate with each other.

More specifically, we consider such networks in a potentially hostile environment, where survivability is paramount. Furthermore, it is assumed that substantial communications capability is required between the network nodes, which may be, for example, command control centers, sensor fields, or storage facilities. The high bit-rate requirements among such nodes, e.g., 10 Gb/s, are best met with a fiber-optic network as opposed to using radio or free-space optics.

One example is a network that has been rapidly deployed for disaster recovery. The network fibers are likely not well protected, such that protection against fiber cuts is needed. Another example is a network near a battlefield or in a small enclave in a war zone. It is assumed that delivering power to the core of the network is not possible; active components, such as transmitters, receivers, and switches, can be supported only at the users at the periphery of the network. The network must be able to survive one or possibly multiple failures. A third example is a network in a military plane or other mobile platform. It may be desirable that the network infrastructure be completely passive to minimize power drain. End-nodes have self-contained powered units (e.g., transmitters, receivers) that are plugged in at various stations on the mobile platform as necessary. As the network may be subject to attack, survivability against failures is critical.

Manuscript received January 24, 2007. Revised July 31, 2007.

J. M. Simmons is with Monarch Network Architects, Holmdel, NJ 07733 USA.

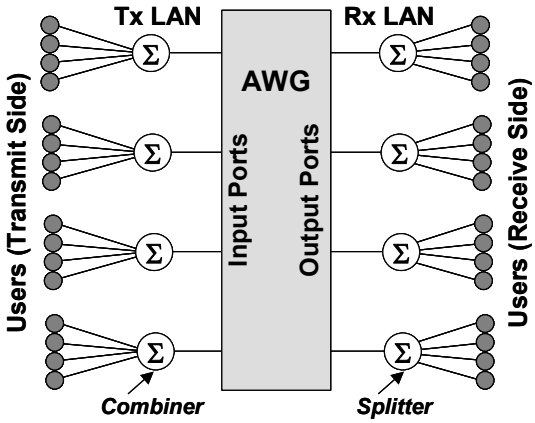

Fig. 1. The AWG star architecture of [3][4]. The transmitter (Tx) output of each user group is passively combined and fed into one input port of the AWG. Each output port of the AWG is broadcast to all receivers (Rx) of the associated user group.

One starting point for such passive networks is the AWG-based star architecture proposed as part of the All-Optical Network program [3][4]. As shown in Fig. 1, users are grouped into LANs, with each transmit LAN feeding into one input port of the AWG, and each receive LAN being fed by one output port of the AWG. As described in [3], this architecture allows any two users to communicate, while being relatively efficient with respect to the number of wavelengths needed. Protection for this architecture, however, was not specifically addressed. A similar architecture was proposed in [5] for airplane-based networks; however, again protection was not specifically addressed.

One strategy for protecting AWG star networks is to deploy parallel copies where the users attach to multiple AWGs, e.g., the dual-star topology of [6]. A disadvantage of this approach is the amount of fiber that needs to be deployed to connect the users to the AWGs. In a hostile or confined environment it may be difficult to route cables such that there are diverse paths from each user to each of the AWGs. Another strategy for providing protection is to attach users to multiple input and

output ports of a single AWG. The multiple paths from user to AWG are routed over diverse fiber cables to provide protection against fiber cuts. As with the dual-star design, finding the necessary diverse paths and routing the required amount of fiber may be difficult. (These methods of protection are also similar in nature to the protection schemes proposed for Gigabit Passive Optical Networks (GPON), where diverse fiber and/or ports are used [7].) To address the cabling issue, [8] proposes a protection strategy for an AWG-based network that is more efficient in its use of fiber. However, the scheme requires active switching in the network, which we assume is not permitted in the relevant environment studied here.

In contrast to the star-based AWG architecture, we propose providing a protected passive environment through the deployment of a small network of AWGs. By arranging the AWGs in survivable topologies and with specific patterns of interconnection between the outputs of one AWG and the inputs of its neighboring AWGs, and by equipping users with tunable transmitters and receivers (or transmitter and receiver arrays), it is possible to provide connectivity between any two users with a high degree of protection. Building a passive network of interconnected AWGs was previously proposed in [9], however, providing protection was not explicitly considered. Cascading a series of AWGs has been proposed, e.g., in [10][11], however the application in these studies is unprotected tree-based access.

Section II reviews the properties of AWGs at a high level; more details can be found in [12]. Section III presents the foundations of designing a distributed network of AWGs. The details of how users interface to such networks are discussed in Section IV. Section V presents several designs that provide protection against one failure. The designs differ in the required end-user equipment and the recovery mechanism. Section VI extends the discussion to architectures that provide protection against two concurrent failures, based on bipartite and cubic arrangements of AWGs. Section VII calculates the expected optical loss of such networks, and confirms the feasibility of the passive infrastructure.

The architectures considered are suitable for tens of users and a small number of AWGs. Larger protected topologies are of course possible from an architectural perspective. For example, a ring with any number of AWGs can be designed that provides protection against any single failure. However, as the number of users and AWGs grows, so does the loss. Thus, practically speaking, the requirement that the network be passive restricts the architecture to relatively small networks.

II. AWG PROPERTIES

AWGs have very structured routing properties. Consider an MxM AWG, and let the input ports and output ports be numbered 0 through M-1. Let λ_0 be the wavelength that when entering on input port 0 exits on output port 0, and assume there is a fixed spacing between wavelengths $\Delta\lambda$. AWGs exhibit cyclic routing properties, where the transfer function repeats after a wavelength interval called the free spectral range (FSR). With an FSR equal to $M \bullet \Delta\lambda$, the AWG can be designed such that, in general, if λ_i enters on input port j, it will exit on output port (i+j) modulo M. In each period of wavelengths, where a period is M consecutive wavelengths, there is exactly one wavelength that goes from each input port to each output port. Thus, to increase the number of simultaneous connections between ports, additional wavelength periods are used.

In order for a totally passive network of AWGs to be feasible from a practical standpoint, there are several required features. First, the AWG needs to be athermal, such that it requires no electrical power for stabilization [13]. Such AWGs are currently commercially available, e.g., [14]. Second, it is desirable that the loss through the AWG be as low as possible because the network will not have amplifiers (except for possibly at the end user locations). There has been much recent work on developing low-loss AWGs, e.g., [15] [16]; loss on the order of less than 4 dB across the band of wavelengths is likely attainable. Third, in the applications discussed here, it is necessary to pass undistorted through a small number of consecutive AWGs (typically less than five); thus, flat, stable passbands are desirable to ensure that a particular wavelength can pass through a cascade of AWGs [17]. Techniques for achieving a flat passband are summarized in [16]. Finally, the number of required input and output ports must be realizable. As the designs discussed here are no larger than 15x15, such AWGs sizes are readily achievable.

The goal of the paper is to lay out the architecture for highly-reliable passive AWG-based networks and emphasize the utility of such an architecture. While the references indicated above provide some measure of confidence that a small network of AWGs is feasible, experimentation is still required to demonstrate this. For example, factors such as crosstalk, dispersion, polarization, and temperature-dependence need to be investigated further.

III. NETWORK OF AWGS

The goal is to provide reliable and substantial communication between users that are geographically distributed over an area, where the network infrastructure is completely passive except for the end-user equipment. The core infrastructure proposed here is based on a network of distributed AWGs, where a subset of the users is attached to each AWG. (The details of how the users connect to the AWGs are discussed in the next section.) The AWGs are arranged in survivable topologies with specific interconnection patterns. We assume that the interconnection configurations are uniform; i.e., each AWG is configured exactly as every other AWG in the network. Clearly other designs are possible where the configurations differ among the AWGs. However, uniform designs are simpler to consider and are sufficient to produce the desired survivability properties. Another advantage of uniform designs is that one can imagine building a 'fiber harness' to manage the fibers coming into and going out of a node; with a uniform design, the same harness can be used at all nodes.

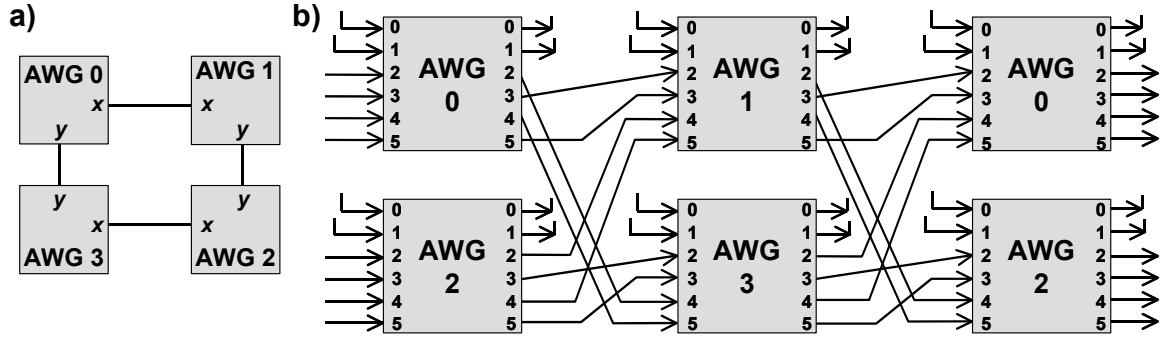

Fig. 2. a) A ring of four AWGs. b) Interconnection pattern of Architecture 1, yielding a diverse path between every pair of routers. For ease of tracing the paths, AWGs 0 and 2 are shown twice in the figure.

A. Network Parameters

Let M represent the number of input ports and the number of output ports of the AWG. Let A be the number of input ports and the number of output ports that are selected for network ingress and egress, respectively (the ingress and egress ports can be selected independently). The remaining M-A input and output ports are used for interconnection. Let each AWG have D directly connected neighbors. Let C represent the number of individual fibers interconnecting any two neighboring AWGs. All ports are utilized; thus: $M = A + D \bullet C$.

A simple example is illustrated in Fig. 2. We assume there are four AWGs, labeled 0 through 3, arranged in a ring topology as shown in Fig. 2(a). As is well known, ring topologies are survivable against all single failures. Each AWG has two directly connected neighbors (i.e., D = 2). The details of how the four AWGs are interconnected are shown in Fig. 2(b). (To simplify illustrating the interconnection pattern, AWGs 0 and 2 are shown twice in the figure.) In the figure, A equals two; i.e., two of the input ports are used for ingress, and two of the output ports are used for egress. The users assigned to a particular AWG gain access to the network though the ingress and egress ports of their respective AWG. In this figure, C is also two, such that there are two (unidirectional) fibers running from an AWG to each neighboring AWG. It is assumed that the C fibers are deployed in a single conduit. Thus, deploying an architecture with C greater than one should require roughly the same effort and only slightly more cost than if C equals one.

Let the two neighbors of each AWG be labeled neighbor X and neighbor Y, where a consistent labeling mechanism must be used. For example, in Fig. 2, the following labeling convention is used: if the AWG IDs are represented using binary notation, then the neighbor that has the right-most bit flipped is labeled neighbor X; neighbor Y is two greater than neighbor X (modulo M). Thus, relative to AWG 1, AWG 0 is neighbor X and AWG 2 is neighbor Y.

In addition to specifying the topology and the parameters A and C, it is necessary to specify the detailed interconnection pattern between neighboring AWGs, as is discussed further in Section III.C. For example, in the figure, output port #3 of each AWG is attached to input port #2 of its neighbor X.

B. Network Paths

After the interconnection pattern is fixed, the supportable paths through the network are determined. As all AWGs are identically configured, it is sufficient to examine the paths that originate on the ingress ports at AWG 0. Due to the periodic routing property of the AWG, it is sufficient to consider only wavelengths 0 through M-1, where M equals 6 in the example. The same paths are generated in all other periods of M wavelengths.

For the configuration shown in Fig. 2(b), the following paths are produced, where the path notation λj: G_1-G_2-...-G_k indicates the j^{th} wavelength is launched from an ingress port of AWG G_1, is routed through AWGs G_2 through G_{k-1}, and exits on an egress port of AWG G_k. (For the paths listed below, G_1 is always AWG 0; the ingress port on AWG 0 is as indicated; the egress port on the destination AWG is not listed.)

Ingress Port 0: λ_0: 0–0; λ_1: 0–0; λ_2: 0–3;
λ_3: 0–1–0; λ_4: 0–3–2; λ_5: 0–1–2–3
Ingress Port 1: λ_0: 0–0; λ_1: 0–3–2–1;
λ_2: 0–1–2; λ_3: 0–3–0; λ_4: 0–1; λ_5: 0–0

Note that there are node/link-diverse paths from AWG 0 to each of the other AWGs. For example, to reach AWG 3, either path 0-3 or path 0-1-2-3 could be used. This allows recovery from any one link failure or any one AWG failure (except the AWGs to which the source and destination users are attached). Some of the paths enter and leave on the same AWG (e.g., wavelength 0 launched from ingress port 0). While these paths may seem extraneous, they are in fact necessary. The paths that start and end on the same AWG are required in order for two users assigned to the same AWG to communicate.

C. Notation

To simplify the discussion of various architectural configurations, the following notation is used to label the input and output ports. I_i is used to indicate the i^{th} ingress port, for $i \in [0, A)$; E_i is used to indicate the i^{th} egress port, for $i \in [0, A)$. A label of X_i on an output port indicates the port is the i^{th} connection between the AWG and its neighboring AWG X, where $i \in [0, C)$. X_i on an input port indicates that

this port is fed by the output port with the same label on the corresponding neighboring AWG (where the input port belongs to the AWG that is neighbor X relative to the AWG with the output port). Similarly, output port Y_i connects to input port Y_i of the AWG that is neighbor Y.

Using this notation, the configuration of Fig. 2(b) is shown in Table 1. We refer to this design as Architecture 1.

TABLE 1. PORT CONFIGURATION FOR ARCHITECTURE 1

Port Number	0	1	2	3	4	5
Input Ports	I_0	I_1	X_0	X_1	Y_0	Y_1
Output Ports	E_0	E_1	Y_0	X_0	Y_1	X_1

IV. ACCESS ARCHITECTURE

Groups of users are assigned to each AWG, where ideally the users are located relatively close to their assigned AWG. We assume each user participates in at most one one-to-one connection at a time, although more general configurations are possible. The outputs of the transmitters of all users assigned to an AWG are passively combined together; the passive combiner, or 'transmit LAN', feeds into one ingress port on the respective AWG. Each AWG egress port feeds into a passive splitter, or 'receive LAN'; the output of the LAN is sent to the receivers of each of the users assigned to the AWG. We assume the transmitters and receivers are tunable over the full range of wavelengths (also, we assume that the transmitters are equipped with an output on-off switch so that the process of tuning the laser does not disrupt other connections). We refer to the fiber that runs between user and AWG as the access fiber. The fiber running between AWGs is referred to as the core fiber.

Consider an architecture where each AWG has two ingress ports and two egress ports. Users potentially need the flexibility to access either of the ingress or egress ports depending on the connections they need to establish. One design is to equip each user with two transmitters and two receivers, as illustrated in Fig. 3(a) (just one AWG is shown in the figure). Two separate transmit LANs and receive LANs are formed, where each transmit LAN is associated with one particular ingress port and each receive LAN is associated with one egress port. The fibers running in the user-to-AWG direction can be deployed in the same conduit as the fibers running from AWG to user. If protection against cuts in the access fiber is desired, then the fiber for one transmit/receive LAN pair should be routed diversely from the fiber for the other transmit/receive LAN pair. (With users relatively close to their associated AWGs, this cabling should not be onerous.)

When a user wishes to establish a new connection, a path is selected that runs between its own AWG and the AWG of the destination user. The selected path determines the required ingress port and wavelength. For example, in Architecture 1, shown in Fig. 2(b), if a user on AWG 0 wishes to communicate with a user on AWG 1 over the path 0-1, then wavelength 4 on ingress port 1 is required. The source user turns on the transmitter corresponding to this ingress port and tunes it to the appropriate wavelength. The destination user tunes the receiver corresponding to the desired egress port to the same wavelength. If the connection is bi-directional, then another path is established in the reverse direction. Unless otherwise indicated, all architectures in this paper support bi-directionally symmetric routing, where the two paths are routed on the same core links, but in opposite directions. The reverse path may, however, be carried on a different wavelength; i.e., the transmitter and receiver of a user participating in a bi-directional connection may be tuned to different wavelengths. (In a small number of paths, the ingress and egress ports used at an AWG may not be symmetric.)

Note that the passive splitter on each receive LAN enables limited multicasting, where one source simultaneously transmits to multiple users on the destination LAN via a single wavelength. However, if the multicast destinations are assigned to different AWGs, then a separate connection needs to be established from the source to each of the destination AWGs, requiring additional transmitters at the source (e.g., the users could be equipped with a transmitter array [18]).

There are various schemes for providing protection using the access architecture of Fig. 3(a). In 1:1 protection, a second path is not established until the first path fails. Establishing the second path requires either re-using the original transmitter and tuning it to a different wavelength, or turning on the second transmitter and tuning it to the appropriate wavelength. Having two transmitters and two receivers potentially also affords the opportunity to establish 1+1 protection, where there is one working path and one active backup path. Users would have a splitter to feed the client signal into both transmitters, and a switch to select the better of the two receiver outputs; these are not shown in the figure.

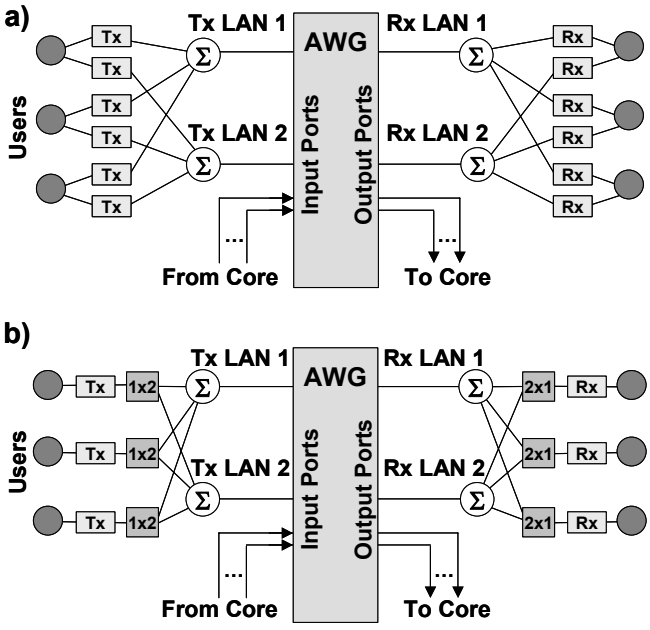

Fig. 3. a) Access architecture where each user is equipped with two transmitters (Tx) and two receivers (Rx). b) Access architecture where each user has a single transmitter and receiver.

Rather than having two transmitters and receivers, another option is for each user to have just one transmitter and a 1x2 switch to select which ingress port is accessed, and one receiver with a 2x1 switch to select the output from the required egress port. This is illustrated in Fig. 3(b). The 1x2 and 2x1 switches are active; however, as was previously mentioned, it is assumed that power is available for end-user equipment. This option is less costly, however, it does not support 1+1 protection, nor does it provide protection against a transmitter/receiver failure.

There are various means of providing the control for the network. For example, there could be one controller for the whole network that is deployed on one of the AWGs (possibly along with some users). Each of the other AWGs establishes a protected bi-directional connection to the controller's AWG with wavelengths that are designated as control wavelengths. The control wavelengths from a particular AWG are shared among the users on the LAN, thereby requiring a Medium Access Control (MAC) protocol such as Aloha. This allows all users to communicate with the controller to establish connections. Furthermore, when a connection fails, the failure is detected by the destination, which then signals the failure to the controller. The controller can then indicate to the source that it should switch to the backup path. (Having the destination signal the failure to the controller is similar in nature to what was demonstrated in [8].) Note that if a protection scheme such as 1+1 is used for the data connection, where there is an active backup path, then there is no need to involve the controller in order to recover from the failure; the client layer at the destination simply chooses the better of the received signals. The controller is informed of the failure after the protection event.

Other control schemes are possible. For example, for extra redundancy, there can be a second controller located on another AWG that becomes active if the primary controller fails. Alternatively, if the users are equipped with transmitter and receiver arrays, then distributed control at the users is possible, eliminating the need for a centralized controller; this type of scheme is described in [18].

V. PROTECTION AGAINST SINGLE FAILURES

Architecture 1, shown in Fig. 2(b), is the simplest four-AWG ring architecture that yields diverse paths between each pair of AWGs. Below, we discuss alternative architectures that yield potentially more useful designs.

A. 1+1 Protection

While Architecture 1 yields a diverse set of paths between every pair of AWGs, full 1+1 protection cannot be supported with this configuration, even if the users have two transmitters and two receivers. This can be seen by noting that the diverse paths to some destinations originate on the same ingress port. For example, paths 0-3 and 0-1-2-3 both originate on ingress port 0. There is only one transmitter per user associated with each ingress port, thus, both of these paths cannot be simultaneously established by the user; i.e., 1:1 protection can be supported, but not 1+1. (This also holds for paths from AWG 0 to AWG 1.)

Full 1+1 protection can be achieved using alternative architectures. Consider a ring of four AWGs where the AWGs have parameters A=2, C=3, M=8, and the interconnection pattern shown in Table 2. (The interconnection pattern becomes more difficult to illustrate as the AWG size increases; thus, for all remaining architectures, we provide only the interconnection table, as opposed to a diagram.) We refer to this design as Architecture 2. Note that the number of interconnecting fibers is greater than in Architecture 1. However, as noted above, all of the fibers between two neighboring AWGs lie in the same conduit. Thus, the effort required to deploy three interconnecting fibers is essentially the same as with two interconnecting fibers. The cost is somewhat higher; however, typically, most of the cost in deploying fiber is incurred when installing the conduit.

The usable network paths provided by Architecture 2, starting from AWG 0, are shown in Table 3. The notation 'i/e(w)' indicates the path is generated by launching wavelength w on ingress port i of AWG 0, where the path exits on egress port e of the destination AWG. In some cases, multiple ingress/egress/wavelength combinations produce the same path. For example, the following combinations generate the path AWG 0 - AWG 3 - AWG 2: wavelength 6 launched on ingress port 0 of AWG 0 and exiting on egress port 0 of AWG 2; wavelength 3 launched on ingress port 1 of AWG 0 and exiting on egress port 0 of AWG 2; and wavelength 4 launched on ingress port 1 of AWG 0 and exiting on egress port 1 of AWG 2.

Some wavelength/ingress-port combinations produce excessively long paths. It is desirable to minimize loss and the number of consecutive AWGs that need to be traversed; thus, any paths that pass through the same AWG twice are considered unusable and are not included in Table 3.

Architecture 2 produces a set of diverse paths between every AWG. Furthermore, the diverse paths utilize diverse ingress and egress ports. Consider a connection between AWG 0 and AWG 1. Path 0-1 starts on ingress port 1 of AWG 0 and ends on egress port 1 of AWG 1, whereas path 0-3-2-1 utilizes ingress port 0 of AWG 0 and egress port 0 of AWG 1. Similar port diversity exists for each of the other connections. Thus, if users are equipped with two transmitters and two receivers, as was shown in Fig. 3(a), 1+1 connections can be established for any source/destination user combination. In addition, if the fibers between user and AWG are diversely routed, then protection against one access fiber failure can be provided.

TABLE 2. PORT CONFIGURATION FOR ARCHITECTURE 2

Port Number	0	1	2	3	4	5	6	7
Input Ports	I_0	I_1	X_1	Y_1	Y_2	X_0	Y_0	X_2
Output Ports	E_0	X_0	X_1	E_1	Y_0	Y_1	Y_2	X_2

TABLE 3. USABLE PATHS FOR ARCHITECTURE 2

Path	Ingress/Egress	Path	Ingress/Egress
0-1	1/1(1)	0-3-2-1	0/0(4)
0-1-2	0/0(2), 0/1(7), 1/1(0)	0-3-2	0/0(6), 1/0(3), 1/1(4)
0-3	0/0(5)	0-1-2-3	1/1(6)
0-0	0/0(0), 0/1(3), 1/0(7), 1/1(2)		

B. Mix of Protected and Unprotected Connections

In addition to providing protected services, assume the network must also support a subset of users that solely want unprotected connections. It is desirable to equip these unprotected users with just a single transmitter and receiver without requiring a 1x2 or 2x1 switch. Architecture 2 can support this configuration. If all users desiring unprotected connections attach to ingress port 0 and egress port 0 it is possible to communicate to any other user. To see this, note that there is at least one path from AWG 0 to each other destination that has the 0/0 port combination. (Alternatively, all unprotected users could attach to ingress port 1 and egress port 1.)

A four-AWG ring architecture that may be better suited to supporting a mix of protected and unprotected connections is configured as follows: A=2, C=3, M=8, with the interconnection pattern of Table 4. We refer to this design as Architecture 3. The usable connection paths are shown in Table 5. One can divide these paths up into one set of diverse paths between each AWG pair and one set of unprotected paths between each AWG pair (the paths designated for unprotected use are underlined in the table). The users solely requiring unprotected connections would attach to ingress port 0 and egress port 0. One drawback of this design as compared to Architecture 2 is that there is not diversity on the ingress and egress ports for all of the protected paths so that protection against a cut in the access fiber is not provided for all connections and 1+1 protection cannot be supported.

TABLE 4. PORT CONFIGURATION FOR ARCHITECTURE 3

Port Number	0	1	2	3	4	5	6	7
Input Ports	I_0	Y_2	I_1	X_0	Y_1	X_1	Y_0	X_2
Output Ports	E_0	X_0	E_1	X_1	Y_0	Y_1	Y_2	X_2

TABLE 5. USABLE PATHS FOR ARCHITECTURE 3

Path	Ingress/Egress	Path	Ingress/Egress
0-1	<u>0/0(3)</u>, 1/1(7)	0-3-2-1	0/1(6)
0-1-2	0/0(7), 1/1(1)	0-3-2	<u>0/0(5)</u>, 1/1(3)
0-3	0/1(4), 1/0(2)	0-1-2-3	<u>0/0(1)</u>, 1/1(5)
0-0	<u>0/0(0)</u>, 0/1(2), 1/0(6), 1/1(0)		

C. More Efficient Use of Wavelengths

Architectures 1, 2 and 3 produce one set of diverse paths between each pair of AWGs, per period of M wavelengths. If the distribution of users and connection load is such that P protected connections need to be simultaneously supported between user groups, then a total of P•M wavelengths are needed.

Consider another four-AWG ring architecture where A=2, C=4, M=10, and the interconnection pattern is as shown in Table 6. We refer to this design as Architecture 4. This arrangement produces the usable network paths shown in Table 7. Two sets of diverse paths between each pair of AWGs are produced in every period. Thus, a total of $\lceil P/2 \rceil \bullet M$ wavelengths are needed to support P simultaneous protected connections between user groups. For P equal to 6, 30 wavelengths are needed. In contrast, Architecture 1 needs 36 wavelengths and Architectures 2 and 3 need 48 wavelengths.

Furthermore, Architecture 4 offers ingress and egress port diversity such that 1+1 protection can be supported. In addition, users that solely require unprotected connections need access only to, for example, ingress port 0 and egress port 0. (For this unprotected traffic arrangement, one path from each AWG will not be bi-directionally symmetric. For example, the path from AWG 0 to AWG 2 follows 0-3-2 but the return path follows 2-1-0.)

TABLE 6. PORT CONFIGURATION FOR ARCHITECTURE 4

Port Number	0	1	2	3	4	5	6	7	8	9
Input Ports	I_0	I_1	Y_3	Y_1	X_0	X_1	X_2	X_3	Y_2	Y_0
Output Ports	E_0	X_0	Y_0	X_1	Y_1	E_1	Y_2	X_2	Y_3	X_3

TABLE 7. USABLE PATHS FOR ARCHITECTURE 4

Path	Ingress/Egress	Path	Ingress/Egress
0-1	0/1(1), 1/1(8)	0-3-2-1	1/0(7), 0/0(2)
0-1-2	0/1(3), 1/1(6)	0-3-2	1/0(5), 0/0(4)
0-3	0/0(8), 1/0(1)	0-1-2-3	1/1(0), 0/1(9)
0-0	0/0(0), 0/1(5), 1/0(9), 1/1(4)		

D. 1:1 Protection Without Requiring Switching

Assume that users have just one transmitter and receiver and that diversity in the access fiber is not required. Furthermore, assume that when a failure occurs, it is desired that the recovery process entail only retuning the transmitter and receiver; configuring the 1x2 and 2x1 switch at the user is not permitted. This requirement could arise, for example, if the switches are slow to reconfigure. These restrictions can be met with a four-AWG ring with the following configuration: A=2, C=4, M=10, and the interconnection pattern of Table 8. We refer to this design as Architecture 5. The resulting usable paths are shown in Table 9.

As an example, consider a 1:1 protected connection from AWG 0 to AWG 1, and assume the initial path is 0-1; the source uses ingress port 0 and the destination uses egress port 0. If this path fails, then the path 0-3-2-1 can be established simply by retuning the transmitter and receiver. (Note that if the user wishes to communicate with a different destination, then a different ingress port may be needed, requiring the user switch to be reconfigured. However, establishment of a new connection is not as time sensitive as recovery from a failure. Thus, the switch reconfiguration time should be sufficient for this function.)

This architecture yields two sets of diverse paths per period, and is thus similar to Architecture 4 in its wavelength efficiency.

TABLE 8. PORT CONFIGURATION FOR ARCHITECTURE 5

Port Number	0	1	2	3	4	5	6	7	8	9
Input Ports	I_0	I_1	X_0	X_3	Y_2	X_2	Y_1	X_1	Y_0	Y_3
Output Ports	E_0	E_1	X_0	X_1	Y_0	X_2	Y_1	X_3	Y_2	Y_3

TABLE 9. USABLE PATHS FOR ARCHITECTURE 5

Path	**Ingress/Egress**	**Path**	**Ingress/Egress**
0-1	0/0(3), 0/0(5), 0/0(7)	0-3-2-1	0/0(4), 0/0(6)
0-1-2	0/0(2), 1/1(2)	0-3-2	0/0(8), 1/1(8)
0-3	1/1(3), 1/1(5), 1/1(7)	0-1-2-3	1/1(4), 1/1(6)
0-0	0/0(0), 0/1(1), 1/0(9), 1/1(0)		

E. Summary of Four-AWG Ring Designs

A number of four-AWG ring designs have been presented, all of which provide at least one set of diverse paths per period. The characteristics of the designs are summarized in Table 10, on page 8. (There is also an architecture with A=2, C=3, M=8 that would have a 'Y' in the final two columns of the table; in the interest of space, the details are not provided.)

VI. PROTECTION AGAINST DOUBLE FAILURES

In this section, we consider architectures that can recover from any combination of two concurrent failures.

A. 3,3 Bipartite Topology

The first architecture studied is the 3,3-bipartite topology composed of six AWGs. The typical depiction of this topology is shown in Fig. 4(a); a depiction that more closely resembles a network deployment is shown in Fig. 4(b). In this topology, each AWG has three neighbors, X, Y, and Z. The labeling convention used is the following: if the AWG IDs are represented in binary, then the neighbor that has the first bit flipped is labeled neighbor X; neighbor Y is two greater than neighbor X, and neighbor Z is two greater than neighbor Y, where the addition is modulo M. (The neighbor assignments in this case are not reciprocal; only the assignments for AWGs 0, 2, and 4 are shown in Fig. 4(a).) The characteristics of the bipartite designs presented below are summarized at the end of this section in Table 17 (see page 9).

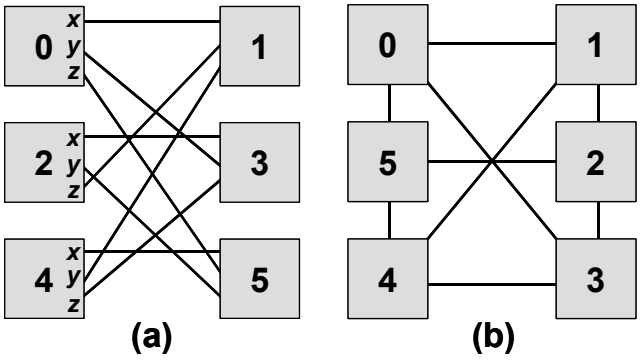

Fig. 4. a) 3,3-bipartite topology. b) Alternate depiction of the topology.

The first two designs discussed below have three ingress and three egress ports. The number of transmitters and receivers required per user depends on the desired level of protection. For example, if 1+2 protection (one working, two active backups) is desired, each user would need to be equipped with three transmitters and three receivers (this is the minimum equipment needed for general 1+2 client layer protection). For 2:1 protection, each user would require just one transmitter and receiver and a 1x3 and 3x1 switch. Another configuration is two transmitters and receivers; one transmitter feeds directly into a transmit LAN and the other transmitter accesses either of the other two transmit LANs via a 1x2 switch. The receiver configuration is similar. This allows two active paths to be initially established; a third path can be established if the first two fail. Any of these schemes allow recovery from two simultaneous failures; the difference is in the time to recover from the failures. A centralized controller can be used to direct the recovery mechanism, as described in Section IV.

B. 1+2 Protection

The first design we consider has the following parameters: A=3, C=3, M=12, and the interconnection pattern of Table 11. We refer to this as Architecture 6. This design produces three diverse paths between every pair of AWGs, as shown in Table 12 (see page 9). Furthermore, it is possible to select diverse paths such that the ingress and egress ports associated with the paths are also diverse, as shown by the underlined entries in the table; this allows 1+2 protection to be implemented, assuming the user has three transmitters and receivers. (If less equipment is available, one of the other protection schemes outlined above can be used.)

TABLE 11. PORT CONFIGURATION FOR ARCHITECTURE 6

Port #	0	1	2	3	4	5	6	7	8	9	10	11
Input Ports	I_0	I_1	I_2	Z_0	Z_1	X_0	Y_0	X_1	Z_2	X_2	Y_1	Y_2
Output Ports	E_0	E_1	E_2	Y_0	Y_1	X_0	Z_0	X_1	Y_2	X_2	Z_1	Z_2

TABLE 10. SUMMARY OF THE FOUR-AWG RING ARCHITECTURES DISCUSSED

Architecture	A	C	M	Produces 2 sets of Diverse Paths per Period	Produces one set of Diverse Paths and one set of Unprotected Paths	Supports 1+1 Protection	Supports 1:1 Protection without Switching	Supports Unprotected Connections via access to one Ingress and one Egress port
1	2	2	6	N	N	N	N	N
2	2	3	8	N	N	Y	N	Y
3	2	3	8	N	Y	N	N	Y
4	2	4	10	Y	N	Y	N	Y
5	2	4	10	Y	N	N	Y	N

C. Mix of Protection Types

Architecture 6 is also well suited for supporting other levels of protection. All users that require protection against only one failure need to access only ingress ports 0 and 2 and egress ports 0 and 2. This allows 1+1 or 1:1 protection. Furthermore, users that require only unprotected connections need to access just ingress port 1 and egress port 1.

D. Protection Without Requiring Switching

Using the configuration A=3, C=3, M=12 did not yield a design where all three paths to a particular destination (in one period) required the same ingress and egress ports. However, it is possible to find a design where two of the three paths to a particular destination enter and leave from the same ingress and egress ports. The interconnection pattern is shown in Table 13; the list of usable network paths is shown in Table 14 (see page 9). This design is referred to as Architecture 7. Switching is not required to recover from the first failure (see the underlined entries in Table 14), but it is required to recover from most of the second failures. For AWG 0 to AWG 2 or AWG 0 to AWG 4, all three paths require only a single ingress and egress port so that switching is not needed to recover from the second failure.

TABLE 13. PORT CONFIGURATION FOR ARCHITECTURE 7

Port #	0	1	2	3	4	5	6	7	8	9	10	11
Input Ports	I_0	I_1	I_2	Y_1	X_1	Y_0	Y_2	Z_0	Z_2	X_2	Z_1	X_0
Output Ports	E_2	Y_0	E_0	Z_0	Y_1	Z_1	Y_2	E_1	Z_2	X_0	X_1	X_2

E. Two-Access Port Designs

Assume that protection against two concurrent failures is required, but that 1+2 protection is not required. Additionally, assume that protection against only one cut in the access fiber is required. These requirements can be satisfied with just two ingress and two egress ports, and two transmitters and receivers per user. The design we consider is: A=2, C=4, M=14, with the interconnection pattern of Table 15. We refer to this design as Architecture 8. As shown in Table 16, this design produces three diverse paths between every pair of AWGs.

This design also provides two options for users requiring protection against just a single failure. 1+1 protection, with diverse ingress and egress ports, can be established. Or, it is possible to operate in a 1:1 mode where switching is not required when one path fails. Finally, unprotected users need to attach to only one ingress and egress port (e.g., ingress port 0 and egress port 0).

TABLE 15. PORT CONFIGURATION FOR ARCHITECTURE 8

Port #	0	1	2	3	4	5	6	7	8	9	10	11	12	13
Input Ports	I_0	Z_0	I_1	X_0	Y_0	X_3	Y_3	Y_1	Z_2	X_2	X_1	Z_3	Z_1	Y_2
Output Ports	E_0	Y_0	E_1	X_0	Z_0	X_1	Z_1	Z_2	Y_1	X_2	X_3	Y_2	Y_3	Z_3

F. Cube Architecture

Another well-known topology that provides protection against two concurrent failures is the 8-node cube, depicted in Fig. 5. As with the 3,3-bipartite design, each AWG has three neighbors. The following labeling mechanism is used: using binary notation, neighbor X is the AWG with the first bit flipped, neighbor Y is the AWG with the second bit flipped, and neighbor Z is the AWG with the third bit flipped.

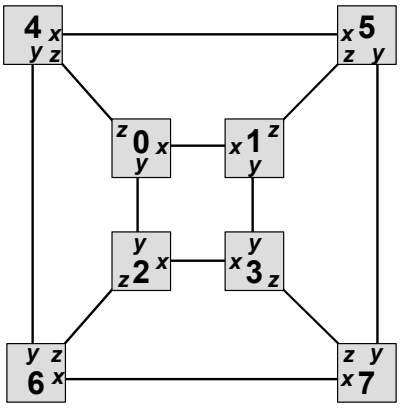

Fig. 5. Cube topology of 8 AWGs

TABLE 12. USABLE PATHS FOR ARCHITECTURE 6

Path	Ingress/Egress	Path	Ingress/Egress	Path	Ingress/Egress
0-1	1/1(6), 0/2(7), 2/0(5)	0-3-2-1	2/2(6)	0-5-4-1	0/0(6)
0-1-2	2/2(3)	0-3-2	0/0(3)	0-5-2	1/1(10)
0-3	1/1(3), 0/2(4), 2/0(2)	0-1-2-3	2/0(7)	0-5-4-3	0/2(11), 2/0(4)
0-1-4	0/0(9)	0-3-4	1/1(2)	0-5-4	2/2(9)
0-5	1/1(9), 0/2(10), 2/0(8)	0-3-2-5	2/0(1), 0/2(8)	0-1-4-5	0/2(5)
0-0	0/0(0), 0/1(1), 0/2(2), 1/0(11), 1/1(0), 1/2(1), 2/0(10), 2/1(11), 2/2(0)				

TABLE 14. USABLE PATHS FOR ARCHITECTURE 7

Path	Ingress/Egress	Path	Ingress/Egress	Path	Ingress/Egress
0-1	1/1(8), 1/1(10), 2/2(8), 0/0(10)	0-3-2-1	2/2(4)	0-5-4-1	1/1(2)
0-1-2	1/0(9)	0-3-2	1/0(5)	0-5-2	1/0(7), 0/2(5)
0-3	0/1(4), 0/2(6)	0-1-4-3	0/1(9)	0-5-2-3	2/2(1)
0-1-4	0/1(11)	0-3-4	0/1(1), 2/0(11)	0-5-4	0/1(3)
0-5	1/0(4), 2/0(6)	0-3-4-5	1/0(3)	0-1-2-5	2/2(7)
0-0	0/0(2), 0/1(7), 0/2(0), 1/0(1), 1/1(6), 1/2(11), 2/0(0), 2/1(5), 2/2(10)				

TABLE 16. USABLE PATHS FOR ARCHITECTURE 8

Path	Ingress/Egress	Path	Ingress/Egress	Path	Ingress/Egress
0-1	1/1(7)	0-3-4-1	0/0(8)	0-5-2-1	0/0(6)
0-1-2	0/0(10)	0-3-2	1/1(13), 0/1(11)	0-5-2	1/1(5)
0-3	1/1(10)	0-1-2-3	0/0(5)	0-5-4-3	0/0(13)
0-1-4	1/1(1), 1/0(3)	0-3-4	1/1(9)	0-5-4	0/0(4)
0-5	1/1(4)	0-3-2-5	0/0(1)	0-1-4-5	0/0(9)
0-0	0/0(0), 0/1(2), 1/0(12), 1/1(0)				

TABLE 17. SUMMARY OF 3,3-BIPARTITE ARCHITECTURES DISCUSSED

Architecture	A	C	M	Supports 1+2 Protection	Supports 1+1 Protection	Supports 1:1 Protection without Switching	Supports Unprotected Connections via just one Ingress and Egress port
6	3	3	12	Y	Y	N	Y
7	3	3	12	N	Y	Y	N
8	2	4	14	N	Y	Y	Y

TABLE 18. PORT CONFIGURATION FOR ARCHITECTURE 9

Port #	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Input Ports	I_0	I_1	I_2	Z_1	Z_0	X_2	Z_3	Y_3	Z_2	X_1	Y_0	X_3	Y_1	X_0	Y_2
Output Ports	E_0	E_1	X_0	E_2	X_1	Y_0	Z_0	Z_1	X_2	Y_1	X_3	Y_2	Z_2	Z_3	Y_3

TABLE 19. USABLE PATHS FOR ARCHITECTURE 9

Path	Ingress/Egress	Path	Ingress/Egress	Path	Ingress/Egress
0-1	0/0(2)	0-2-3-1	2/1(9)	0-4-5-1	0/1(13)
0-2	0/0(5)	0-1-3-2	1/2(9)	0-4-6-2	0/0(12)
0-1-3	2/1(2)	0-2-3	2/1(7)	0-4-5-7-3	1/0(11)
0-4	1/2(12), 2/2(10)	0-2-6-4	2/0(3)	0-1-5-4	0/0(4)
0-1-5	1/0(7)	0-4-5	0/2(7)	0-2-3-7-5	0/2(11)
0-2-6	0/2(14)	0-4-6	1/2(6)	0-1-5-7-6 0-4-5-7-6	1/1(3) 2/1(11)
0-1-5-7	0/0(8)	0-2-3-7	2/1(12)	0-4-6-7 0-4-5-7	1/2(5) 2/1(4)
0-0	0/0(0), 0/1(1), 0/2(3), 1/0(14), 1/1(0), 1/2(2), 2/0(13), 2/1(14), 2/2(1)				

One cubic design that produces three diverse paths to each node is: A=3, C=4, M=15, with the interconnection pattern of Table 18. We refer to this design as Architecture 9. The usable paths are shown in Table 19. (There are two 'extra' paths produced, 0-4-5-7-6 and 0-4-5-7, which are not needed to achieve 3-way path diversity. Not all paths are bi-directionally symmetric; e.g., 0-1-5-7 vs. 7-3-1-0; however, a bi-directional connection between 0 and 7 is still resilient against two core failures.)

While this design provides protection against two failures in the core, it protects against only a single failure in the access fiber; i.e., there is only two-way diversity in the ingress and egress ports, not three-way. This two-way diversity is shown by the underlined entries in Table 19. A search was not performed for cubic designs with other survivability features.

VII. OPTICAL LOSS

The network infrastructure is required to be passive, thus, the total optical loss on any path must be low enough that amplifiers in the core fiber are not needed. Systems have been demonstrated that tolerate up to 39 dB loss for 10 Gb/s connections with no amplifiers required along the fiber [19]; thus, this loss figure is used here as a feasibility guideline for a passive system. If the actual tolerable losses are smaller, then either the geographic extent or the number of users would need to be reduced. To analyze the expected optical loss that arises from the designs presented here, the following assumptions are made: Fiber Loss: 0.2 dB/km; AWG Loss: 4 dB; distance from users to AWG: <1 km. Assume there are 24 users in the network.

For the four-AWG ring designs, assume the distance between AWGs is 5 km (airplane distances are clearly much shorter). The worst-case path traverses the access links at the source and destination, three core links, and four AWGs. The loss of this path is: $(1+3\cdot5+1)\cdot0.2 + 4\cdot4 \approx 19$ dB. Assume there are six users per AWG, such that there is a ~8 dB splitting loss on the transmit LAN and on the receive LAN. Additionally, assume there is another 2 dB of miscellaneous loss to account for factors such as excess loss in the splitters and couplers. The total loss is on the order of 37 dB, which, as noted above, should be feasible for a passive system.

For the 3,3-bipartite designs, assume the distance between AWGs is 5 km, except for the diagonal links shown in Fig. 4(b), which have a distance of roughly 11 km. The worst-case path traverses the access links at the source and destination, three core links (one diagonal link and two shorter links), and four AWGs. The loss of this path is: $(1+11+2\cdot5+1)\cdot0.2 + 4\cdot4 \approx 21$ dB. Assume there are four users per AWG, such that there is a 6 dB splitting loss on the transmit LAN and on the receive LAN. Adding another 2 dB of miscellaneous loss yields a total loss on the order of 35 dB.

For the cube design, assume the distance between AWGs on the inner square of Fig. 5 is 5 km and the distance between AWGs on the outer square is 10 km. The diagonal links have a distance of roughly 3.5 km. The worst-case path traverses the access links, four core links (two outer links, one inner link, and one diagonal link) and five AWGs. The loss of this path is: $(1+2\cdot10+5+3.5+1)\cdot0.2 + 5\cdot4 \approx 26$ dB. Assume there are three users per AWG, such that there is a ~5 dB splitting loss on the transmit LAN and on the receive LAN. Adding another 2 dB of miscellaneous loss yields a total loss on the order of 38 dB. A geographically smaller network would allow more margin. For example, in an airplane environment, the total loss would be on the order of 33 dB. Note that the worst-case path is needed only when there are two failures. With zero or one failure, only three links and four AWGs are traversed, and the worst-case loss is ~6 dB less.

VIII. CONCLUSION

A variety of passive protected networks can be constructed by interconnecting a small number of AWGs. The designs presented protect against either a single failure or a double failure, where the failure could be any core link, or any AWG

except for the source and destination AWG. Some architectures also provide protection against single or double failures along the access fiber. While some amount of cabling is required to deploy these architectures, it is simplified by the fact that all fiber between two neighboring AWGs lies in the same conduit. The cabling is further simplified by the fact that the interconnection pattern of each AWG is identical, allowing a single fiber harness configuration to be used at each AWG.

In extremely hostile environments, it may be desirable to provide protection against more than two simultaneous failures. One architecture that increases the survivability is a double-ring design, where two independent four-AWG rings are deployed. Users requiring extreme protection would connect to one AWG on each of the two rings. This would provide protection against three simultaneous fiber cuts. Furthermore, it provides protection if one of the AWGs to which the user is attached fails.

Alternatively, one could deploy topologies that provide more protection. For example, 4,4-bipartite topologies (with 8 AWGs) and degree-4 hypercubes (with 16 AWGs) provide protection against up to three failures. Users could be connected to two AWGs to gain protection against failure of their primary AWG. While highly redundant, these topologies may be difficult to realize in practice due to the large amount of required cabling and the potentially high loss.

A future paper will cover the strategies used for designing general protected architectures and will cover designs based on alternative wavelength routing devices.

REFERENCES

[1] N. Frigo, et al., "Approaches to multiple service delivery over passive optical networks," *OFC'98*, San Jose, CA, Feb. 22-27, 1998, pp. 404-405.

[2] M. Maier and M. Reisslein, "AWG-based metro WDM networking," *IEEE Communications Magazine*, November 2004, pp. S19-S26.

[3] S. Alexander, et al., "A precompetitive consortium on wide-band All-Optical Networks," *Journal of Lightwave Technology*, Vol. 11, No. 5/6, May/June, 1993, pp. 714-735.

[4] R. Barry and P. Humblet, "On the number of wavelengths and switches in All-Optical Networks," *IEEE Transactions on Communications*, Vol. 42, No. 2/3/4, Feb./Mar./Apr., 1994, pp. 583-591.

[5] I-S. Joe and O. Solgaard, "Scalable optical switch fabric for avionic networks," *Avionics Fiber-Optics and Photonics*, Minneapolis, MN, Sept. 20-22, 2005, pp. 19-20.

[6] Y. Sakai, et al., "Management system for full-mesh WDM AWG-star network," *ECOC*, Amsterdam, Sept. 30-Oct. 4, 2001, pp. 264-265.

[7] ITU-T Recommendation G.984.1, "Gigabit-capable Passive Optical Networks (GPON): General characteristics," March 2003.

[8] A. Hill, et al., "Multiple-star wavelength-router network and its protection strategy," *IEEE Journal of Selected Areas in Communications*, Vol. 16, No. 7, September, 1998, pp. 1134-1145.

[9] F. Banerjee, and B. Mukherjee, "Passive optical network architecture based on wavelength grating routers," *IEEE Journal of Selected Areas in Communications*, Vol. 16, No. 7, September, 1998, pp. 1040-1050.

[10] M. Parker, F. Farjady, and S. Walker, "Wavelength-tolerant optical access architectures featuring N-dimensional addressing and cascaded arrayed waveguide gratings," *Journal of Lightwave Technology*, Vol. 16, No. 12, December 1998, pp. 2296-2302.

[11] G. Maier, et al., "Design and cost performance of the multistage WDM-PON access networks," *Journal of Lightwave Technology*, Vol. 18, No. 2, February 2000, pp. 125-142.

[12] C. Dragone, "An NxN optical multiplexer using a planar arrangement of two star couplers," *IEEE Photonic Technology Letters*, vol. 3, No. 9, September 1991, pp. 812-815.

[13] J. Hasegawa and K. Nara, "Low loss ~1.4 dB 200GHz-16ch athermal AWG compact module for metro/access network," *OFC'05*, Anaheim, CA, Mar. 6-11, 2005, OTuD5.

[14] NeoPhotonics, Athermal AWGs, Data Sheet.

[15] C.R. Doerr, L.W. Stulz, and R. Pafchek, "Compact and low-loss integrated box-like passband multiplexer," *IEEE Photonics Technology Letters*, Vol. 15, No. 7, July 2003, pp. 918-920.

[16] C. Dragone, "Low-loss wavelength routers for WDM optical networks and high-capacity IP routers," *Journal of Lightwave Technology*, Vol. 23, No. 1, January 2005, pp. 66-79.

[17] T. Otani, et al., "Cascadability of passband-flattened arrayed waveguide-grating filters in WDM optical networks," *IEEE Photonics Technology Letters*, Vol. 11, No. 11, November 1999, pp. 1414-1416.

[18] C. Fan, S. Adams, and M. Reisslein, "The FTL–FRL AWG Network: A Practical Single-Hop Metro WDM Network for Efficient Uni- and Multicasting," *Journal of Lightwave Technology*, Vol. 23, No. 3, March 2005, pp. 937-954.

[19] G. Lenz and M. Blasio, "Raising dispersion tolerance to increase distance," *Lightwave*, December 2003, p. 13.