Analysis of Internal ROADM Protection

Jane M. Simmons
Monarch Network Architects
Holmdel, NJ USA

Adel A. M. Saleh
University of California - Santa Barbara
Santa Barbara, CA USA

***Abstract*—Reconfigurable optical add/drop multiplexers (ROADMs) are a key component in most carrier networks. One common ROADM design utilizes a set of wavelength-selective switches (WSSs) in a broadcast-and-select architecture. All traffic at the node passes through at least one WSS. If a WSS fails, the affected traffic can be routed over spare network capacity, similar to a link failure. Alternatively, WSS redundancy can be provided internal to the ROADM. We discuss the advantages of internal ROADM protection, and analyze two architectures for providing protection for the WSSs.**

Keywords—broadcast-and-select, internal protection, reconfigurability, ROADM, wavelength-selective switch, WSS

I. Introduction

Reconfigurable optical add/drop multiplexers (ROADMs) are a key component of carrier transport networks. Each network node is typically equipped with a ROADM and all traffic at the node, whether it be *transiting* traffic or *add/drop* traffic, passes through the ROADM. ROADMs provide wavelength-level switching from any input network fiber to any output network fiber; enable optical bypass, where transiting traffic remains in the optical domain to reduce cost and power consumption; and provide the gateway to the optical network for client services (e.g., IP) that originate/terminate at the node. They are a critical enabler of reconfigurable networks; their inherent agility has become essential as applications such as cloud computing and software defined networking (SDN) have begun to burgeon.

Given their central role in the transport layer, a ROADM failure is at least as serious as a link failure (possibly worse), and needs to be addressed expeditiously. One option is to make use of protection capacity deployed in the network to re-route traffic that has been brought down by the ROADM failure. A second option is to provide redundancy within the ROADM itself so that the affected network traffic can be restored without being re-routed. Internal restoration offers several advantages, as discussed in Section II.

We focus on providing internal protection for the most common ROADM architecture: broadcast-and-select (B&S) based on a small set of wavelength-selective switches (WSSs) [1]. WSSs can direct a wavelength from any input port to any output port. In the B&S architecture, the WSS is likely the most vulnerable of the ROADM components; thus, the analysis here specifically focuses on protection for a WSS failure.

Although WSS-based ROADMs were first proposed in the early 2000s, there is little published work analyzing their protection. One early paper addressing ROADM protection is [2], which includes comparisons of network availability for systems with and without internal ROADM protection. Using the network to provide protection via traffic rerouting is not considered. Also, the baseline architecture of [2] deviates somewhat from the common commercial ROADM configuration, as discussed further below. Reference [3] analyzes the availability of various ROADM architectures, including several WSS-based architectures; however, internal protection is not considered. The availability of 'programmable' ROADMs is the focus of [4] and [5]. In this proposed configuration, an optical backplane is used to arbitrarily interconnect the various internal components of the ROADM. This flexibility enables internal protection for elements such as the WSSs. The programmable ROADM is a relatively new proposal and is not discussed further here.

The paper is organized as follows. The baseline unprotected broadcast-and-select WSS-based ROADM architecture is presented in Section II. Section III presents two architectures for providing internal protection for a failed WSS. These architectures are compared in section IV with respect to availability, loss, cost, and failure coverage.

II. Baseline B&S ROADM Architecture

A. WSS-Based Architecture

There are several variations of the B&S ROADM architecture, one of which is illustrated in Fig. 1. The figure depicts a degree-two node (i.e., two network fiber-pairs); however, the discussion applies to a node of any degree. There are two add/drop ports shown. The particular architecture shown in Fig. 1 is *directionless*, where a signal on any of the add (drop) ports can be directed to (received from) any of the output (input) network fibers.

On the input side, all traffic passes through a 1×4 passive splitter that *broadcasts* each signal to each one of the output directions. Each output direction is equipped with a 4×1 WSS to *select* which of the signals should continue on to a network fiber or a drop port. The ROADM is often augmented by an edge cross-connect (XC), as shown, which allows the client services to access any of the add/drop ports.

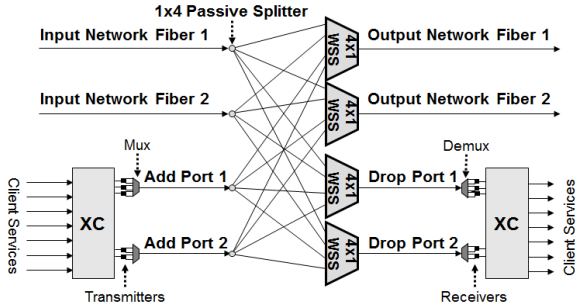

Fig. 1. Baseline unprotected B&S WSS-based ROADM.

(Our discussion is focused on protecting the output WSSs. However, the mux/demux component of the ROADM may be a WSS as well. Recovery from a failure of a mux/demux WSS can be provided by the edge XC.)

B. Network Protection for a WSS Failure

First, consider a failure of a WSS that feeds an output network fiber. The ramifications are similar to a failure of the fiber itself. Any traffic that was routed on that fiber, whether it be transiting traffic or sourced traffic, is brought down. Protection can be provided via the network, where the affected traffic is re-routed around the failure using spare capacity.

If the affected traffic is sourced at the node, then it may need to be launched on a different wavelength in order to avoid wavelength contention on the recovery path. The add ports carry WDM signals; thus, the new wavelength must also be free on the add port. If it is not free on the original add port, then the edge XC can be used to move the affected connection to another add port. It may be desirable to deploy additional add/drop ports to reduce such wavelength contention issues.

Next, consider a failure of a WSS that feeds a drop port. Such a WSS processes traffic that terminates at, or is regenerated at, the node. Recovery requires that the affected traffic be received on another drop port; the directionless property of the ROADM allows the affected traffic to continue to be received from the same input network fiber. However, moving traffic to a different drop port may lead to wavelength contention on the new port. If a free wavelength cannot be found on both the new drop port and the existing input network fiber, then the connection may need to be re-routed using a different input network fiber.

C. Internal Protection for a WSS Failure

Rather than recovering from a WSS failure by re-routing traffic in the network, another option is to include a spare WSS and confine the recovery operation to within the ROADM. There are several advantages to this approach.

First, spare network capacity is not needed for recovery, allowing this capacity to be utilized to recover from other concurrent failures. It also increases the availability of the spare capacity for low-priority pre-emptible traffic. To gauge the potential benefit, we studied several backbone networks and assumed that link failures result from fiber cuts, line-amplifier failures, and WSS failures. If network protection is utilized for a WSS failure, then the link associated with that WSS can be considered failed until the WSS is replaced. With internal ROADM protection, the switchover to the spare WSS occurs almost immediately and the link downtime is negligible. Overall, the study showed that internal ROADM protection reduced the expected amount of time that a network suffers from one or more link failures by up to 10%.

Second, and perhaps more importantly, utilizing a spare WSS for recovery maintains the paths/wavelengths of the affected traffic, avoiding issues with wavelength contention. In a heavily loaded network, relying on traffic re-routing for protection against a WSS failure may not always be feasible because of wavelength conflicts. This may force a subset of the affected connections to remain down until the failed WSS is replaced.

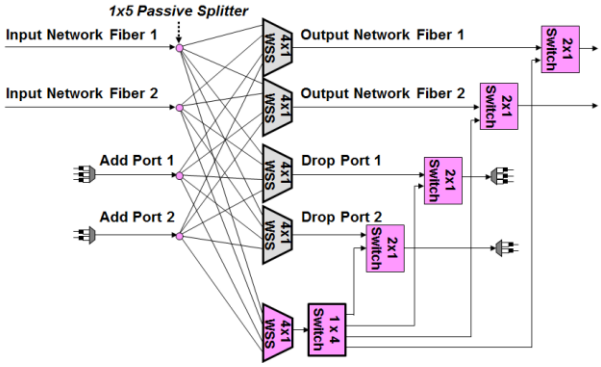

Fig. 2. Protected Architecture 1.

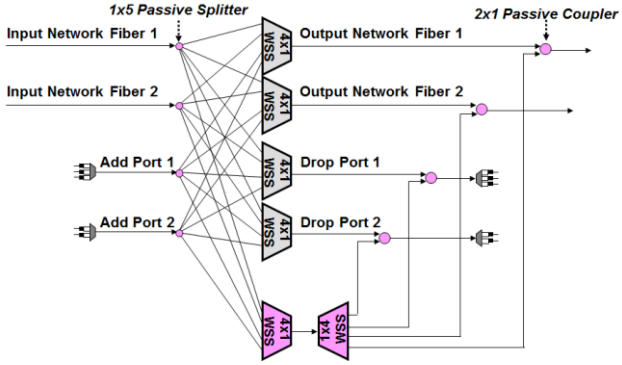

Fig. 3. Protected Architecture 2.

Third, if network protection is accomplished via a shared mesh scheme, then internal recovery via a spare WSS is likely to be significantly faster; e.g., connections may be restored in roughly tens of msec as opposed to hundreds of msec.

The disadvantages of internal protection are the cost of the spare equipment and the slightly more complex architecture of the ROADM.

III. Internal ROADM Protection Architectures

Two architectures that provide internal WSS protection are described next. (If the ROADM utilizes output nodal amps, then a spare amp can be included with the spare WSS.)

A. Architecture 1

One architecture that incorporates a spare WSS is shown in Fig. 2 (the edge XC is not shown). As compared to Fig. 1, the changes are: the broadcast splitters are 1×5 instead of 1×4, a spare 4×1 WSS is added, the spare WSS is followed by a 1×4 switch, and a 2×1 switch is added at each of the outputs. The spare WSS takes on the wavelength routing configuration of whichever WSS fails; the resulting signal is directed via the 1×4 switch to the output corresponding to the WSS failure. Each 2×1 switch selects a signal from either the output WSS (if it has not failed) or from the spare WSS. (This architecture is an extension of the internal protection scheme proposed for early-generation broadcast-and-select ROADMs based on dynamic spectral equalizers in the 2000 timeframe [6].)

The protection architecture of Fig. 1 is similar to that of [2]. However, [2] deployed the WSSs on the input side, which can be problematic due to leakage issues during the switching process [1]. Additionally, the protection architecture of [2] requires an additional bank of 1×2 switches.

B. Architecture 2

The second ROADM protection architecture, shown in Fig. 3, was presented in [5]. The broadcast splitters are 1×5 instead of 1×4, a spare 4×1 WSS is added, the spare WSS is followed by a 1×4 WSS, and a 2×1 coupler is added at each of the outputs. The spare WSS takes on the configuration of the failed WSS and directs its signal to the proper output via the 1×4 WSS. The 2×1 coupler combines signals from the corresponding output WSS and the spare WSS.

Under most conditions, the design assumes that only one signal is present at the 2×1 coupler at a given time. However, if a WSS has failed, then that WSS may not be able to block all wavelengths from passing through to the WSS output ports; i.e., signals may 'leak' through from the failed WSS and interfere with the signal coming from the spare WSS. Thus, a practical implementation of Architecture 2 may replace the 2×1 couplers with 2×1 switches.

IV. Architectural Comparisons

This section compares the two protection architectures with respect to availability, loss, cost, and failure coverage. The assumptions regarding the major components are shown in Table I. The results are summarized in Table II.

A. Through-Path Availability (Row 1 of Table II)

We focus on the availability of a *through* path for a ROADM with a total of *N* output network fibers and drop ports (*N* equals 4 in Fig. 1). With no internal protection, the availability is dominated by the WSS on the output fiber; thus, the availability is ~ A_W. (It is assumed that the splitter has an availability of ~100%.)

1) Architecture 1

The j^{th} *through* path is 'UP' if:

(j^{th} 2×1 switch is UP) AND ((j^{th} WSS is UP) OR ((j^{th} WSS is DOWN) AND (Spare WSS is UP) AND (1×N switch is UP) AND (no other path is using the spare WSS)))

For N=4, the availability can be expressed as:

$$A_S \bullet [A_w + (1\text{-}A_w) \bullet (A_w) \bullet (A_S) \bullet (A_w^3 + [3\,A_w^2(1\text{-}A_w)]/2 + [3\,A_w(1\text{-}A_w)^2]/3 + (1\text{-}A_w)^3/4)]$$

For arbitrary *N*, this generalizes to:

$$A_S \bullet [A_w + (A_w \bullet A_S)\,(1 - A_w^N)/N] \approx A_S = 0.9999992$$

The availability is dominated by the 2×1 switch.

2) Architecture 2

The j^{th} *through* path is 'UP' if:

(j^{th} 2×1 coupler is UP) AND ((j^{th} WSS is UP) OR ((j^{th} WSS is DOWN) AND (Both the spare WSS and the 1×N WSS are UP) AND (no other path is using the spare WSS)))

For N=4, the availability can be expressed as:

$$A_C \bullet [A_w + (1\text{-}A_w) \bullet (A_w^2) \bullet (A_w^3 + [3\,A_w^2(1\text{-}A_w)]/2 + [3\,A_w(1\text{-}A_w)^2]/3 + (1\text{-}A_w)^3/4)]$$

For arbitrary *N*, the general expression for availability is then:

$$A_C \bullet [A_w + (A_w^2)\,(1 - A_w^N)/N] \approx A_C \approx 1.0$$

TABLE I. Component Availability and Loss Assumptions

	WSS	Small Switch	1×N Coupler/Splitter
Availability	$A_W = 0.99998$	$A_S = 0.9999992$	$A_C \approx 1.0$
Nominal Loss	6.5 dB	1 dB	$10 \log_{10} N$

TABLE II. Results

		No Internal Protection	Arch. 1	Arch. 2 (with 2×1 coupler)	Arch. 2 (with 2×1 switch)
Availability of a Through Path		~ A_W	~ A_S	~ A_C	~ A_S
Through-Path Loss (for N=4)	**Not Using Spare WSS**	12.5 dB	14.5 dB	16.5 dB	14.5 dB
	Using Spare WSS		15.5 dB	23 dB	21 dB
Approx. Additional Cost		-	Cost of one WSS	Cost of two WSSs	Cost of two WSSs
Failure Coverage		-	1 WSS failure	Fully recover from one WSS failure; Partially recover from multiple WSS failures	

Thus, with a coupler on the outputs, Architecture 2 provides higher availability than Architecture 1. However, if 2×1 switches replace the 2×1 couplers, as discussed in the previous section, then the availability is approximately A_S in both architectures.

B. Through-Path Loss (Rows 2 and 3 of Table II)

With no internal protection, the 1×N splitter and the outgoing WSS are the main contributors to the nominal through-path loss of 12.5 dB (for $N = 4$).

1) Architecture 1

A through-path that does not utilize the spare WSS suffers the additional loss of the 1×(N+1) splitter and the loss of the 2×1 output switch. If the through-path makes use of the spare WSS, then the loss also includes that of the 1×N switch. As indicated in Table II, these additional losses are on the order of 2 to 3 dB. The extra loss can be compensated for by increased gain of the nodal optical amplifier that is typically present. There should not be a major degradation of the achievable optical reach. (The optical reach is the distance an optical signal can travel in the network before the signal quality degrades to a level that necessitates regeneration.)

2) Architecture 2

A through-path that does not utilize the spare WSS suffers the additional loss of the 1×(N+1) splitter and the loss of the 2×1 coupler (or 2×1 switch). A through-path utilizing the spare WSS also incurs the loss of the 1×N WSS. As shown in Table II, there is an appreciable increase in path loss, as compared to Architecture 1, when the spare WSS is utilized. Connections must continue to remain feasible when shunted through the spare WSS; thus, the relatively large loss of the internal ROADM protect path must be considered when determining the optical reach. While increased nodal

amplification can be used to compensate for the increased loss, it is likely that the optical reach of Architecture 2 is shorter than that of Architecture 1. This typically would result in more regeneration, especially in a backbone network.

C. Cost (Row 4 of Table II)

Virtually all of the additional cost in both architectures is represented by the added WSSs. Thus, Architecture 2, with two additional WSSs, adds roughly twice as much to the cost as does Architecture 1, which has only one additional WSS.

D. Failure Coverage (Row 5 of Table II)

1) Architecture 1

Architecture 1 is capable of providing protection for all traffic affected by a single WSS failure.

2) Architecture 2

Architecture 2 similarly can provide protection for all traffic affected by a single WSS failure. Additionally, it is capable of providing protection for *a subset* of the affected traffic if more than one WSS fails. The spare N×1 WSS can select wavelengths from any of the failed WSSs, and direct these wavelengths to the 1×N WSS. This second WSS then directs the wavelengths to the proper output ports. The limitation is that the restored wavelengths must be unique to avoid wavelength contention on the single line between the N×1 and the 1×N WSSs. For example, one could choose to restore the highest priority affected traffic, subject to the unique-wavelength restriction. The traffic on the remaining affected wavelengths would need to be re-routed in the network in order to be restored.

As an extension of the multiple-WSS-failure scenario, consider the case where one or more of the WSS failures is only a *partial* failure. With a partial failure, some of the wavelengths can still be directed from a WSS input port to the proper output port (e.g., a MEMS-based WSS where only a subset of the mirror elements have failed). Thus, these wavelengths do not need to be shunted through the spare WSS. Additionally, assume that Architecture 2 employs 2×1 couplers at the outputs rather than 2×1 switches. Then wavelengths arrive at an output port via both its associated partially-failed WSS (i.e., the wavelengths unaffected by the failure) *and* via the spare WSS (i.e., some or all of the wavelengths affected by the failure). The 2×1 coupler combines the two wavelength streams to form the output WDM signal. (The two wavelength streams carry disjoint wavelengths such that there are no wavelength contention issues at the output coupler.) Note that this protection mode is not supported if Architecture 2 is equipped with 2×1 switches at the outputs, nor is it supported by Architecture 1. The same restriction as noted above, namely that the wavelengths restored via the spare WSS must be unique, still holds. Thus, it is still possible that some of the traffic brought down by the failed/partially-failed WSSs cannot be restored internal to the ROADM.

If it were desirable to restore *all* traffic taken down by two concurrent WSS failures, then the spare N×1 and 1×N WSSs could be replaced by N×2 and 2×N WSSs, respectively.

V. Conclusion

Protection from ROADM failures warrants the attention of service providers. ROADMs are a fundamental element of optical networks, especially as paradigms such as SDN engender greater network configurability. Furthermore, the ROADM failure rate, at least in the near term, is likely to increase as more complex components, such as bandwidth-variable WSSs, are introduced.

The advantages of internal ROADM protection were discussed and two potential architectures were analyzed. Architecture 1 has lower loss and cost; Architecture 2 provides more protection if multiple WSSs fail. The discussion extends to the route-and-select ROADM architecture as well.

References

[1] J. M. Simmons, Optical Network Design and Planning, 2nd ed., Springer, 2014, Chapter 2.

[2] A. Morea and I.B. Heard, "Availability of translucent networks based on WSS nodes, comparison with opaque networks," *ICTON 2006*, Nottingham, UK, June 18-22, 2006, Paper Mo.P.11.

[3] M. Dzanko, et al., "Analytical and simulation availability models of ROADM architectures," *International Conference on Telecommunications*, Zagreb, Croatia, June 26-28, 2013.

[4] M. Dzanko, et al., "Availability analysis of optical cross-connect implemented by architecture on demand," *ICTON 2012*, Coventry, England, July 2-5, 2012, Paper Tu.C2.5.

[5] M. Dzanko, M. Furdek, G. Zervas, and D. Simeonidou, "Evaluating availability of optical networks based on self-healing network function programmable ROADMs," *J. Opt. Commun. Netw.*, vol. 6, no. 11, Nov. 2014, pp. 974-987.

[6] M. Reardon, "Corvis Upgrades Optical Switch," *Lightreading*, March 20, 2002, Available Online: http://www.lightreading.com/ethernet-ip/corvis-upgrades-optical-switch/d/d-id/579268.