A Closer Look at ROADM Contention

Jane M. Simmons

***Abstract*— Reconfigurable optical add/drop multiplexers (ROADMs) provide the wavelength-switching capability in the optical layer of most transport networks. As the need for greater network configurability grows, ROADMs continue to evolve to provide greater flexibility. More specifically, ROADMs that provide colorless, directionless, gridless, and contentionless operation are on the roadmap of many equipment vendors and service providers. The first three of these properties are easy to define. ROADM contention, however, can take on many forms, which has led to some misconceptions regarding the contentionless property. We examine ROADM contention in more detail in order to provide greater clarity regarding contentionless ROADMs.**

I. Introduction

Reconfigurable optical add/drop multiplexers (ROADMs) provide wavelength-granularity switching and the ability to add/drop individual wavelengths to/from a wavelength-division multiplexed (WDM) signal. Their ability to perform these operations in the optical domain has led to a dramatic reduction in the amount of electronics required at network nodes, thereby providing benefits in cost, power, space, and reliability. First deployed in carrier networks in the 2000 time-frame, ROADMs have undergone numerous architectural improvements over the past 15 years, as described in [1]-[4]. Most recently, the focus has been on adding operational flexibility to ROADMs, to better enable configurable (as opposed to quasi-static) networks.

Four properties have received the most attention with respect to flexible ROADMs: colorless, directionless, gridless, and contentionless. The first three of these properties are relatively simple to describe. *Colorless* indicates that any slot of the ROADM can accommodate a transponder of any wavelength. (A transponder is a transmit/receive card. A slot is the physical position on a ROADM shelf where the transponder card is inserted.) This is especially beneficial when using tunable transponders as it allows a transponder to be tuned to a different wavelength without having to manually move the transponder to a different ROADM slot. It also eliminates operational errors where a transponder is inserted in the 'wrong' slot. *Directionless* refers to the ability of a transponder to access any of the network fibers that enter/exit a ROADM. Thus, a transponder that initially launches a connection on the eastbound fiber may later launch a connection on the westbound fiber, without requiring any manual intervention. This is useful for optical-layer restoration and for dynamic networking in general. *Gridless* indicates that the ROADM filters can accommodate a range of wavelength spacings and modulation formats. This is needed to support, for example, flexible transponders and elastic optical networks [5]. Gridless ROADMs are also referred to as bandwidth-variable ROADMs.

ROADM contention can take on many forms, leading to some confusion as to what constitutes a contentionless ROADM. We define the contentionless property as the ability of the ROADM to support any connection that can be supported by the current network configuration. Here we focus on wavelength contention (contention can also arise, for example, due to limits on the amount of add/drop at the ROADM [6]). In that context, contentionless operation implies that if a wavelength is available to be used on a network fiber, it will not be blocked from being used by the ROADM. Wavelength contention in a ROADM can largely be avoided through the use of effective algorithms (e.g., by treating the ROADM add/drop ports as 'network links' in standard routing and wavelength-assignment (RWA) algorithms). Simulation studies that demonstrate the benefits of intelligent algorithms with regard to ROADM wavelength contention can be found in [7]. Nevertheless, ROADM wavelength contention still occurs when the network is heavily loaded and/or when the network is highly dynamic. There are also restoration schemes that have a tendency to engender wavelength contention, as discussed in Section III.C below. Thus, there is continuing effort to develop cost-effective contentionless ROADMs.

There are three canonical ROADM architectures (see [1]-[3] for more details of the architectures). For illustrative purposes regarding wavelength contention, we use the broadcast-and-select (B&S) architecture with optical splitters on the input ports and small wavelength-selective switches (WSSs) on the output ports. A directionless configuration is assumed, with the add/drop ports treated similarly to the network-fiber ports. (The add/drop ports are the interface to the 'clients' in the electronic layer; the network-fiber ports are the interface to the fiber pairs that run between neighboring nodes.) Most of the contention scenarios enumerated below apply to the route-and-select (R&S) ROADM architecture as well, where WSSs are utilized on both the input and output ports. The third ROADM architecture, the wavelength-selective architecture, is based on a large centralized all-optical switch [8]. It is inherently contentionless and will not be discussed further here. It has found limited deployment due to concerns about its reliability and scalability (although [8] presents a design that addresses this latter concern).

In Section II, we illustrate the ROADM contention that may occur when the number of add/drop ports is less than

Manuscript received August 29, 2016.

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

the number of network fibers. One of the more common misconceptions is that if the number of add/drop ports equals the number of network fibers, then wavelength contention will not occur within the ROADM. Sections III and IV provide examples of where this is not true. Another assumption that has appeared regarding ROADMs is that the combination of the colorless and directionless properties implies contentionless operation. The examples of Sections II through IV show this to be an erroneous assumption as well.

It is important to understand the circumstances under which contention occurs in order to design contentionless ROADMs. There have been claims of ROADM architectures being contentionless when in fact they have been shown to exhibit contention under particular scenarios (an example of this is explored in [9]). As we illustrate the various forms of ROADM contention, we also include a discussion of how the contention can be avoided. The various architectures for avoiding contention have ramifications for cost and ancillary supported functionality. Section V addresses how contentionless operation is currently being designed into ROADMs.

II. Wavelength Contention Due To A Limited Number Of Add/Drop Ports

The first example of wavelength contention is shown using the B&S ROADM of Fig. 1. The inputs to the ROADM are shown at the left of the figure; the outputs are shown on the right. The ROADM is assumed to be deployed at a degree-three node (the nodal degree indicates the number of input/output network fiber-pairs at the node).

Note that the add/drop ports carry WDM signals. While the number of add/drop ports typically equals the number of network fibers, in this particular example there are only two add/drop ports rather than three. This design decision may be made to reduce cost: only five WSSs are needed rather than six, and the WSSs are of size 5×1 as opposed to 6×1.

Assume that three connections arrive to this ROADM, where all three connections need to be dropped to the electronic layer (e.g., an IP router at the node). In addition, assume that all three connections have been routed using λ1 (clearly they must be routed on three different network fibers). Multiple connections on the same wavelength cannot be carried by a given add or drop port due to the resulting contention in the WDM signal. Thus, at most two of the λ1 connections can be accommodated at this node; the remaining connection will be blocked because of ROADM contention. The origin of the contention, i.e., too few add/drop ports, is readily seen in this simple example.

It should also be noted that the B&S ROADM in Fig. 1 is directionless. A transmitter/receiver on any add/drop port can access any of the network fibers. If the mux/demux operation on the add/drop ports (as represented by the trapezoids in Fig. 1) is provided by, for example, WSSs, then the ROADM is also colorless. Thus, as illustrated by this example, and as stated earlier, the combination of the directionless and colorless properties does not imply the contentionless property. (In fact, in the earliest B&S ROADMs of the 2000 time-frame, the add/drop ports were simple taps off of each of the network fibers, resulting in a *non-directionless* architecture [1][10]. This early architecture could be considered contentionless, though limited in flexibility, because an add/drop port carried the same WDM signal as its associated network fiber. Thus, if a wavelength was free to be used on a network fiber, it was free to be used on the add/drop port.)

To address the contention described above with respect to Fig. 1, one can add a third add/drop port to the design such that the number of add/drop ports equals the number of network fiber ports. This would allow three λ1 connections to be added/dropped. However, note that the optical splitter on each ROADM input enables multicast in the B&S architecture. For example, a signal that enters the ROADM from the East network fiber could be multicast to drop ports 1 and 2. If we modify the above scenario such that one of the three λ1 connections is required to be multicast to two different drop ports, then having three drop ports is insufficient to prevent contention. Again, one of the desired λ1 connections will be blocked. More generally, if one considers multicast drops, then even with a large number of drop ports, wavelength contention is not theoretically eliminated; although, if one has enough ports, then such contention would be unlikely to occur in practice.

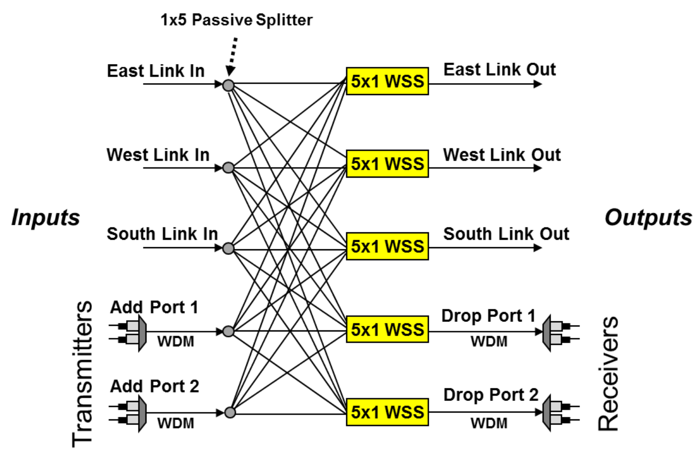

Figure 1. A broadcast & select ROADM with three network input/output fiber pairs and two add/drop ports. The ROADM inputs are on the left; the outputs are on the right. If it is desired that three connections, all routed on λ1, be sourced/terminated at this node, then one of these connections will be blocked due to wavelength contention on the add/drop ports.

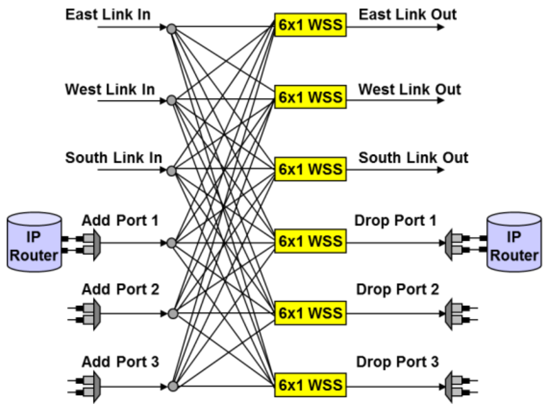

Figure 2. An IP router, which interfaces to add/drop port 1, wants to establish a connection on the East and West links. If λ1 is the only available wavelength on both of these links, then one of the connections will be blocked due to wavelength contention on the add/drop port.

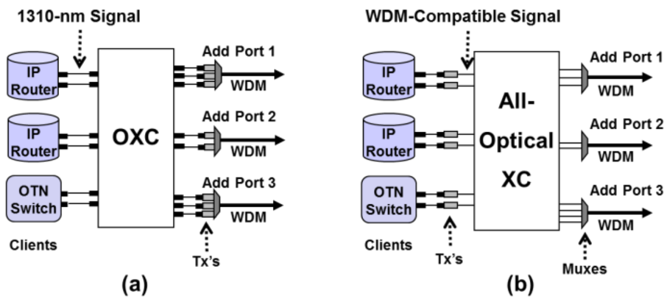

Figure 3. Adding an edge cross-connect to the ROADM. Only the add-side of the architecture, with the multiplexers (muxes) and the transmitter (Tx) portion of the transponders, is shown. (a) The edge optical cross-connect (OXC) is placed between the clients and the transmitters, allowing the clients to direct traffic to any transmitter. This OXC can be electronic, as shown, or all-optical. (b) The edge cross-connect (XC) is placed between the WDM-compatible output of the transmitters and the WDM multiplexer. The transmitters can access any add port. In this architecture, the XC must be all-optical, as shown.

III. Wavelength Contention Due To Limited Edge Configurability

A. *Client Access*

An example of wavelength contention that cannot be ameliorated by an increase in the number of add/drop ports is shown in Fig. 2. For simplicity, a single network client is shown; i.e., an IP router that is connected to add/drop port 1. Assume that the IP router wants to establish two connections, one on the East link and one on the West link. Assume that λ1 is the only available wavelength on these two links. Then one of the connections will be blocked because two λ1 connections cannot be carried on add/drop port 1.

In this scenario, the wavelength contention arises from the lack of edge flexibility. The IP router interfaces to just a single add/drop port. In order to eliminate this type of contention, one solution is to have the router directly interface to transponders on all of the add/drop ports. Alternatively, an edge switch can be added to allow the IP router to access all of the add/drop ports. Adding an edge switch can also support other desirable functions, such as path restoration, dynamic connection establishment, and 1:N protection of the transponders.

There are two possible locations for inserting the edge switch, as shown in Fig. 3 (only the add side is shown in this figure). One option is to place the switch between the clients and the transponders, as shown in Fig. 3a. Alternatively, it could be placed between the transponders and the multiplexers/demultiplexers, as in Fig. 3b. Either option provides the edge configurability necessary to avoid the wavelength contention described above. However, there are subtle implications in the option that is chosen. First, there are likely to be differences as to how many transponders are needed to support a particular traffic pattern. The option that results in fewer required transponders depends on the traffic, as examined in more detail in [1]. Second, in Fig. 3a, the edge switch operates on a standard 1310-nm signal, thus allowing this switch to be electronic. In Fig. 3b, the edge switch operates on the WDM-compatible outputs of the transponders and must be an all-optical (i.e., photonic) switch. Because of concern regarding optical-amplifier transients, an all-optical switch will likely need to be switched more slowly than its electronic counterpart. Thus, from the perspective of restoration and rapid dynamic networking, the architecture of Fig. 3a may be preferred.

Adding an edge switch is shown to be effective in reducing ROADM wavelength contention in the studies of [11][12], where it is assumed that the edge switch is positioned as in Fig. 3a. Furthermore, these studies demonstrate that deploying multiple smaller edge switches is almost as effective at reducing blocking as deploying a single large edge switch. The modular architecture is advantageous in that it provides a 'pay-as-you-grow' cost structure and it avoids having one large cross-connect as a single point of failure.

B. *Regeneration*

Another example of wavelength contention that can be attributed to a lack of edge configurability arises when regenerator cards are used for regeneration. Regenerator cards are functionally similar to two back-to-back transponders, with the short-reach interfaces of the transponders eliminated to save cost. The signal to be regenerated enters the node on network fiber *i* using wavelength λi and is dropped to the regenerator card. The signal is then added by the regenerator card using wavelength λj and exits the node on network fiber *j*. Typically, a regenerator card permits λi to be different from λj. We assume that a regenerator card is inserted into a shelf slot associated with a single add/drop port. With this assumption, contention in the ROADM will occur when the desired λi and λj from the network-fiber perspective are not both available on that add/drop port.

There are a few solutions that can be considered to address this type of contention. First, it can be avoided (or at least minimized) if the ROADM architecture allows the two 'halves' of the regenerator card to access different add/drop ports, e.g., via an optical backplane. This is a form of edge configurability, where the regenerator card would be connected to the add/drop port(s) with λi and λj available. Second, a separate edge cross-connect can be added as discussed above in relation to Fig. 3. However, note that the configuration of Fig. 3a is not compatible with regenerator cards because there is no short-reach interface (i.e., no 1310-nm signal) to feed into the edge switch. In contrast, the configuration of Fig. 3b can be utilized with regenerator cards, as shown in Fig. 4a. This architecture utilizes four input/output ports on the edge cross-connect per regenerator card.

Alternatively, two discreet back-to-back transponders can be used for regeneration as opposed to the single regenerator card (thereby losing the cost benefit of the regenerator card). This architecture can be used in combination with the edge switch configuration of Fig. 3a, as shown in Fig. 4b. This provides flexibility with regard to utilizing the desired add/drop ports to avoid contention, as well as flexibility with regard to interconnecting transponders. Two input/output ports on the edge switch are utilized per regeneration with this architecture. There must be available transponders deployed on each add/drop port to ensure that the two interconnected transponders can support λi and λj without encountering wavelength contention.

If the edge switch utilizes the configuration of Fig. 3b, then the regenerator architecture with back-to-back transponders is as shown in Fig. 4c. Note the similarity of Fig. 4c to Fig. 4a. The switch only provides flexibility with respect to the add/drop ports that are utilized. The transponders would need to be interconnected via patch cords, which is problematic if the transponders are not tunable (many wavelength combinations could be required). Furthermore, four input/output ports on the edge switch are utilized per regeneration with this configuration. Thus, this configuration addresses contention, but with less efficiency and somewhat less flexibility than the configuration of Fig. 4b.

C. *Shared Restoration*

As mentioned above, optical-amplifier transients can be problematic in an all-optical network when switching wavelengths in the optical domain, or when turning up or bringing down wavelengths on a fiber. To avoid transients when restoring traffic after a failure, various protection schemes based on pre-lit wavelengths have been proposed, where a set of pre-lit segments are stitched together at the time of failure in order to form a restoration path [13]. Because the segments are pre-lit, no change in power is experienced on the fiber links.

Using an edge switch to stitch together the pre-lit segments is straightforward, as described in [13]. However, if an edge switch is not deployed, then the segments can be concatenated by reconfiguring the WSSs within the ROADM. Consider the example shown in Fig. 5. It is assumed that there are two pre-lit segments in this example, one carried on λ1 and one on λ2. The node shown in Fig. 5 serves as an endpoint of both pre-lit segments. Figure 5a depicts the baseline configuration, prior to any failure, with λ1 dropping to port 1 and λ2 dropping to port 3. Note that for each pre-lit segment, the receiver (on the right) is tied to its corresponding transmitter (on the left).

Assume that a failure occurs such that the two pre-lit segments need to be concatenated to form part of the restoration path. This can be accomplished by reconfiguring the WSSs on drop ports 1 and 3, with the resulting desired configuration shown in Fig. 5b. The wavelengths on the East and West network fiber pairs remain unchanged, thereby avoiding issues with optical transients. However, this concatenation can only be accomplished if λ1 is available on drop port 3 and λ2 is available on drop port 1. Thus, wavelength contention in the ROADM can potentially block the setup of the desired restoration path.

If an edge switch is used to perform the concatenation, the wavelengths on the network fibers and the add/drop ports are unchanged by the operation, such that wavelength contention in the ROADM is not an issue. This is true for either configuration of Fig. 3. (With the configuration of Fig. 3b, the pre-lit-segment receivers need to be tied to the corresponding transmitters via a patch cord, similar to what is shown in Fig. 5. Additionally, the receivers utilized for the pre-lit segments need to be tunable.)

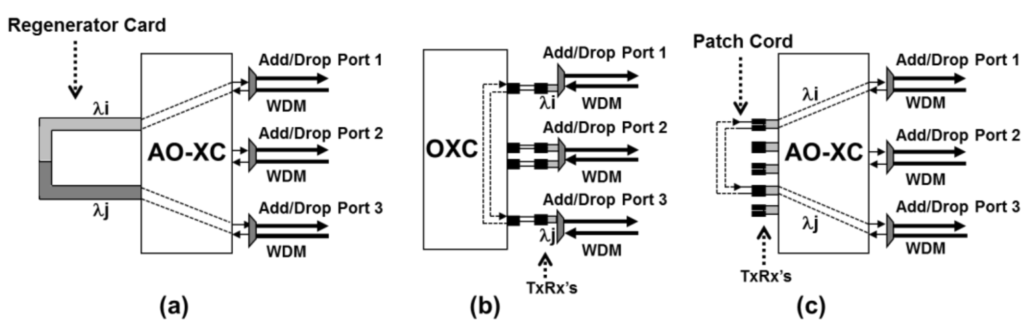

Figure 4. Regeneration at a ROADM equipped with an edge switch. Both add and drop ports are shown. (a) A regenerator card in combination with an all-optical cross-connect (AO-XC). Four input/output ports on the AO-XC are utilized per regeneration. (b) Regeneration is achieved using two back-to-back transponders (TxRx's) interconnected by the OXC. Two input/output ports on the OXC are utilized per regeneration. (c) Two back-to-back transponders in combination with an AO-XC. Four input/output ports on the AO-XC are utilized per regeneration.

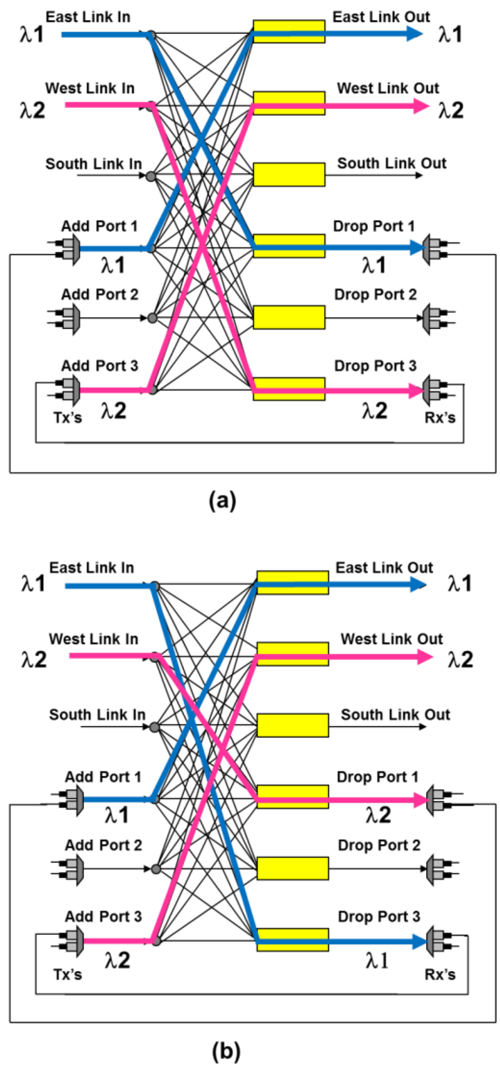

Figure 5. Shared restoration where pre-lit segments are concatenated to form the restoration path. Two pre-lit segments are shown, one on λ1 and one on λ2. (a) The configuration prior to failure. (b) The desired configuration after a failure occurs, where the λ1 segment must be concatenated to the λ2 segment. To support this operation, λ2 must be available on drop port 1 and λ1 must be available on drop port 3.

D. Sliceable Transponders

ROADM wavelength contention potentially has implications for the deployment of sliceable transponders, a technology that has been proposed to enable cost-effective elastic optical networks (EONs) [14][15]. A single physical sliceable transponder supports multiple 'virtual transponders', each of which may be assigned a flexible amount of spectrum (e.g., by adjusting the assignment of subcarriers). The 'optical flows' represented by each of the virtual transponders can be routed independently in the network.

If sliceable transponders are used in an architecture where all optical flows that correspond to a given sliceable transponder are multiplexed together, then 'spectral contention' limits the flexibility of assigning spectrum to the flows. (Spectral contention in EONs is analogous to wavelength contention in conventional grid-based networks.) More specifically, the spectrum that is assigned to the set of virtual transponders associated with a given sliceable transponder cannot overlap. This is illustrated in Fig. 6a, where the edge cross-connect is located before the transponders on the add side of the ROADM (this corresponds to the architecture of Fig. 3a).

In contrast, the architecture shown in Fig. 6b potentially permits more flexibility in assigning spectrum to the flows (this corresponds to the architecture of Fig. 3b). Assuming that there are multiple interfaces between the sliceable transponder and the all-optical switch (as is shown), then the optical flows of a given sliceable transponder can be directed to different ports on the all-optical switch and ultimately to different add ports on the ROADM. This allows the optical flows of one sliceable transponder to be assigned overlapping portions of the spectrum (assuming the technology of the transponder supports this). Reference [15] includes simulations that probe the reduction in blocking probability that can be achieved with this additional flexibility in spectrum assignment.

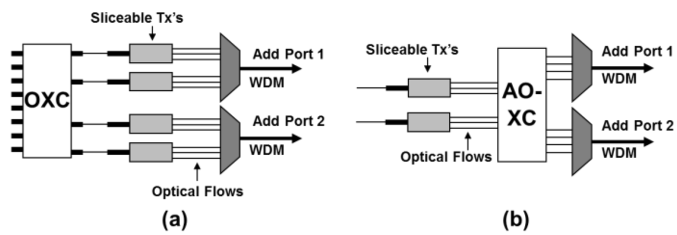

Figure 6. The impact of the edge cross-connect location on the spectrum assignment process when used with sliceable transponders. Only the add-side, with the transmitter (Tx) portion of the sliceable transponder, is shown. (a) With the OXC located before the transmitters, the optical flows emanating from one sliceable transponder are muxed together onto one add port of the ROADM. The spectrum assigned to the flows must be non-overlapping, otherwise wavelength contention occurs on the add port. (b) With the AO-XC located after the transmitters, the optical flows of one sliceable transponder may be directed to different add ports of the ROADM. This potentially allows the spectrum assigned to the optical flows of one sliceable transponder to overlap.

IV. Wavelength Contention Due To Limited Pre-Deployed Equipment

Another example of wavelength contention in a ROADM is portrayed in Fig. 7. In this scenario, we assume that dynamic services are supported by the network, such that there must be equipment pre-deployed in the network; i.e., connections must be established rapidly with no time for a 'truck-roll' to install the required equipment. We assume that the pre-deployed transponders are as shown in the figure. Some of the transponders are assumed to be utilized for existing connections, as indicated. There are a total of two available transponders, one each on add/drop ports 1 and 2.

Assume that a dynamic connection needs to be established on the South link, and assume that the only free wavelength on this fiber is λ1. However, it is assumed that existing connections are already utilizing λ1 on add/drop ports 1 and 2. λ1 is available on add/drop port 3; however, there is no available transponder on this add/drop port. Thus, the new desired connection would be blocked. Note that adding an edge switch using the architecture of Fig. 3a does not address this type of wavelength contention, whereas the architecture of Fig. 3b does. Alternatively, pre-deploying more transponders can minimize the frequency of such a scenario occurring, where the service provider must determine the right tradeoff between the cost of pre-deploying extra transponders and the probability of blocking.

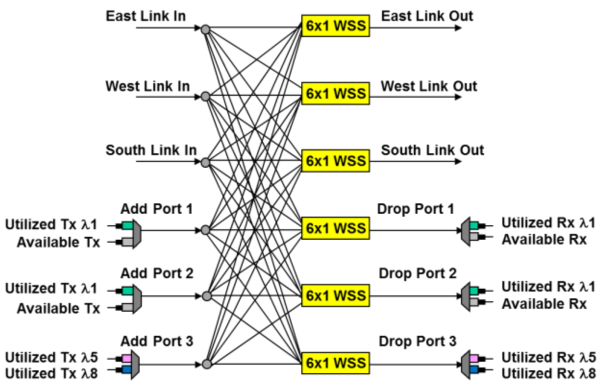

Figure 7. Assume that a new connection is desired on λ1. The only available transponders (TxRx's) are located on add/drop ports 1 and 2. However, the only add/drop port with λ1 available is port 3. Thus, the new connection would be blocked.

V. M×N WSS-Based Contentionless ROADM

As has been pointed out in the previous sections, adding edge configurability can avoid all, or at least most, instances of wavelength contention. The contentionless ROADMs that have been commercially developed thus far combine this edge configurability with the add/drop structure, where a single add/drop port equipped with an M×N WSS [16] replaces the individual add/drop ports shown in previous figures. An M×N WSS allows any of the M inputs to be directed to any of the N outputs of the WSS. M represents the number of transponders at the node, and N represents the number of network fiber pairs. See Fig. 8, where N equals 3.

On the input side, there is a one-to-one correspondence between the N input network fibers and the N inputs of the drop-side N×M WSS. Thus, if a wavelength is available on an input network fiber, it is available on the corresponding drop line as well. On the output side, there *potentially* could be pre-lit wavelengths on the outputs of the add-side M×N WSS that are not passed through to any of the output network fibers (perhaps in support of some type of dynamic networking scheme where pre-lit wavelengths are kept in a stand-by mode until needed to rapidly establish a new connection). Thus, *conceivably* there could be wavelength contention on the add lines. However, under normal network operation, this would not occur. Thus, this architecture is typically considered to be contentionless.

Of course, the M×N WSS itself must be contentionless in order for the ROADM to be contentionless. For example, constructing an M×N WSS from the combination of an M×1 WSS feeding into a 1×N WSS would lead to wavelength contention due to the single line connecting the constituent WSSs.

The M×N WSS architecture is similar, though not identical, to the edge-switch architecture of Fig. 3b. The M×N WSS essentially subsumes the AO-XC and mux/demux functionality. One aspect where the architectures potentially differ is with respect to multicast. If the M×N WSS is not capable of multicast, then multicast drop, where an incoming signal is dropped to multiple clients, is not supported in the architecture of Fig. 8. Multicast add, where an added signal is sent to multiple output network fibers would not be supported either. In contrast, the architecture of Fig. 3b in combination with a B&S ROADM supports both of these multicast operations, though its support of multicast drop is limited to the scenario where the clients are located on different drop ports.

The ROADM of Fig. 8 is also colorless and directionless. Such a ROADM is often denoted as being *CDC* (colorless, directionless, contentionless). The gridless property depends on the flexibility of the filter technology used in the WSSs.

An alternative architecture that is likely more practical replaces the M×N WSS with an M×N multicast switch (MCS) [17]. An MCS provides benefits in cost and space as compared to the WSS; it also supports multicast drop. However, it requires the receivers to be capable of selecting the desired wavelength from a WDM signal, and it typically incurs greater loss.

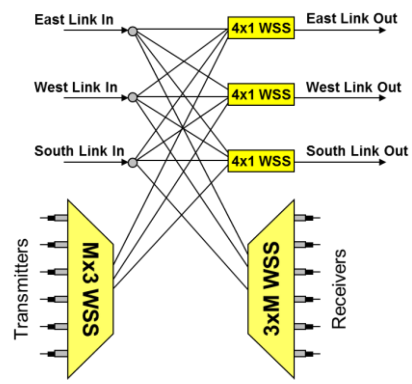

Figure 8. An M×N WSS is deployed to provide contentionless operation. Alternatively, an M×N MCS can be used.

VI. Conclusion

We have described a variety of situations where wavelength contention can occur within a ROADM, along with architectures to avoid the contention. No single architecture is perfect in addressing all of the potential scenarios that may arise. Furthermore, some of the solutions may add substantial cost to the ROADM. As noted in the introduction, many of the contention scenarios can be minimized through the use of effective algorithms. The contentionless property is thus somewhat less crucial as compared to the colorless, directionless, and gridless properties. Nevertheless, simplifying network operation is desirable if it can be accomplished at reasonable cost, such that further research into architecting cost-effective contentionless ROADMs is warranted.

References

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

[2] S. Gringeri, B. Basch, V. Shukla, R. Egorov, and T. J. Xia, "Flexible architectures for optical transport nodes and networks," *IEEE Commun. Mag.*, vol. 48, no. 7, Jul. 2010, pp. 40-50.

[3] B. C. Collings, "Advanced ROADM technologies and architectures," *Proc. Optical Fiber Communication Conf.*, Los Angeles, CA, 2015, Paper Tu3D.3.

[4] S. Perrin, "Next-generation ROADM architectures and benefits," Heavy Reading White Paper, March 2015, Available at: www.fujitsu.com/us/Images/Fujitsu-NG-ROADM.pdf, Accessed on December 8, 2016.

[5] S. Poole, S. Frisken, M. Roelens, and C. Cameron, "Bandwidth-flexible ROADMs as network elements," in *Proc. Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Los Angeles, CA, 2011, Paper OTuE1.

[6] F. Naruse, Y. Yamada, H. Hasegawa, and K. Sato, "Evaluations of OXC hardware scale and network resource requirements of different optical path add/drop ratio restriction schemes," *J. Opt. Commun. and Netw.*, vol. 4, no. 11, Nov. 2012, pp. B26-B34.

[7] P. Palacharla, X. Wang, I. Kim, D. Bihon, M. D. Feuer, and S. L. Woodward, "Blocking performance in dynamic optical networks based on colorless, non-directional ROADMs," *Proc. Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Los Angeles, CA, 2011, Paper JWA8.

[8] J. M. Simmons and A. A. M. Saleh, "Wavelength-selective CDC-ROADM designs using reduced-sized optical cross-connects," *Photon. Technol. Lett.*, vol. 27, no. 20, Oct. 15, 2015, pp. 2174–2177.

[9] T. Zami, P. Jenneve, and H. Bissessur, "Fair comparison of the contentionless property in OXC," *Proc. Asia Commun. and Photonics Conf.*, Hong Kong, 2015, Paper AM3G.3.

[10] A. Tzanakaki, I. Zacharopoulos, and I. Tomkos, "Optical add/drop multiplexers and optical cross-connects for wavelength routed networks," *Proc. International Conference on Transparent Optical Networks,* Warsaw, Poland, 2003, Paper Mo.B2.4, Fig. 4b.

[11] T. Zami, "Contention simulation within dynamic, colorless and unidirectional/multidirectional optical cross-connects," *Proc. European Conference and Exhibition on Optical Communication*, Geneva, 2011, Paper We.8.K.4.

[12] I. Kim, P. Palacharla, X. Wang, D. Bihon, M. D. Feuer, and S. L. Woodward, "Performance of colorless, non-directional ROADMs with modular client-side fiber cross-connects," *Proc. Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Los Angeles, CA, 2012, Paper NM3F.7.

[13] J. M. Simmons, "Cost vs. capacity tradeoff with shared mesh protection in optical-bypass-enabled backbone networks," *Proc. Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Anaheim, CA, 2007, Paper NThC2.

[14] O. Gerstel, "Flexible use of spectrum and photonic grooming," *Proc. International Conf. on Photonics in Switching*, Monterey, CA, 2010, Paper PMD3.

[15] M. Dallaglio, T. Zami, N. Sambo, A. Giorgetti, A. Pagano, E. Riccardi, and P. Castoldi, "Add and drop architectures for multicarrier transponders in EONs,"*J. Opt. Commun. Netw.*, vol. 8, no. 7, Jul. 2016, pp. A12-A22.

[16] B. C. Collings, "Wavelength selectable switches and future photonic network applications," *Proc. International Conference on Photonics in Switching*, Pisa, Italy, 2009, pp. 52-55.

[17] W. I. Way, "Optimum architecture for M×N multicast switch-based colorless, directionless, contentionless, and flexible-grid ROADM," *Proc. Optical Fiber Communication/National Fiber Optic Engineers Conf.*, Los Angeles, CA, 2012, Paper NW3F.5.

Jane M. Simmons [F'11] has been involved in the research and development of optical networks for 20 years. She currently is the founder of Monarch Network Architects, which provides optical network architectural services and design tools. She served as the subject matter expert on optical networking for DARPA for six years, and made significant technical contributions to DARPA's *CORONET* program on dynamic, highly-reliable networks. From 1999 to 2002, she was the Executive Engineer of Network Architecture and later the Chief Network Architect of Corvis Corp. While at Corvis, she performed the network design and the link engineering for the first commercially deployed all-optical backbone network. Through pioneering algorithmic and architectural optimizations, she played a significant role in the adoption of all-optical networking in telecommunications networks. Prior to Corvis, Dr. Simmons worked at Bell Labs/AT&T Labs Research, where she conducted research on backbone, regional, and broadband access networks. She received a B.S., Summa Cum Laude, from Princeton University, and S.M. and Ph.D. degrees from MIT, all in Electrical Engineering.

Dr. Simmons is a Fellow of the IEEE and is a member of Phi Beta Kappa, Tau Beta Pi, and Sigma Xi honor societies. She has served on the technical program committee of several conferences, including *OFC* and *ICC*. She was the OFC Networks Subcommittee Chair in 2003, and was a member of the OFC Steering Committee for six years. From 2004 to 2009, she was an Associate Editor of the *IEEE J-SAC Optical Communications and Networking Series*. From 2009 to 2012, she was an Associate Editor for the *Journal of Optical Communications and Networking*, and is currently on its steering committee. She teaches a course on optical network design at *OFC*, and is the author of the textbook *Optical Network Design and Planning*, now in its second edition.