

power from different orthogonal modes [32]), so all of the field and power of these other three modes goes through the lower three waveguides WA2 – WA4. Hence it does not matter for those modes if we interrupt the top waveguide for add-drop functions. By similar arguments, we can see that this process will also reconstruct the add-drop mode at the output with the signal from transmitter T1; in the case of the add-drop mode, it does not matter for the form of the output if the field in waveguide WA1(O) comes directly from waveguide WA1(I) or if it comes from the transmitter T1.




Fig. 2. Device as in Fig. 1(a) but with added dummy MZIs (the devices in grey, without specific labels), set in their “bar” state and with a standard phase shift, for greater equality of loss and path length.

So far, we have presumed loss-less components, which leads to the unitarity of the matrices U and V . The arguments will remain valid if there is equal overall additional loss for all waveguide paths from the input waveguides to the center waveguides and from the center waveguides to the output waveguides; in such a case, each of U and V is a unitary matrix within a single multiplying constant and the necessary orthogonality properties of functions are retained even if there is overall loss in the system. Since the MZIs may have some loss associated with them, one possible strategy is to ensure that all paths pass through equal numbers of MZIs by adding dummy “bar-state” MZIs [26]. For example, as shown in Fig. 2, by completing rectangular blocks of MZIs by, on the left, adding 3 in WI1, 2 each in WI2 and WA4, and 1 each in WI3 and WA3, and similarly for a rectangular block on the right, we would ensure all left-to-right paths went through 8 MZIs, 4 on each side. Another approach to this device that uses fewer MZIs overall and does not require dummy devices would be to exploit the binary tree architecture in [26] with mostly-transparent detectors and with waveguides similarly connecting between the corresponding points on the two sides, though that approach can require crossing waveguides.

If the overall form of all the modes is to be retained in passing through the device, it is important that, other than the phase differences deliberately imposed with the MZIs, the phase delays in different paths should be essentially equal, at least modulo 2π . Otherwise, the device will affect the form of the other modes (i.e., those orthogonal to the add-drop mode). They will still be orthogonal on leaving the device (and will still be orthogonal to the add-drop mode), but they will be changed by such undesired phase differences in the paths. So that the behavior can be substantially independent of wavelength, it is also desirable that the total path lengths of each waveguide path from input to output are substantially equal overall. Otherwise, as wavelength is changed there is relative phase change between paths of different lengths; that relative phase change will prevent one SRADM device setting from working with multiple different wavelengths in extracting the add-drop spatial mode, and will upset the correct reconstruction of the other spatial modes at different wavelengths. The schemes in Fig. 1 have such substantial equality of paths, and adding the dummy MZIs as in Fig. 2 retains and possibly enhances such equality. Note in Figs. 1(b) and 1(c) that waveguide lengths are added on the two inner paths to and from the grating couplers so as to equalize the overall waveguide lengths. The extent to which a device like this is truly independent of wavelength depends on the details of the dispersion in the materials and response in the Mach-Zehnder devices, which in turn depend on the precise fabrication details. To the extent that waveguide

Mach-Zehnder interferometers with equal arms can be regarded as non-dispersive in other telecommunications applications (for example, over the telecommunications C band), they can likely similarly be reasonably non-dispersive here.

Whether we can use the same device settings for different wavelengths also depends on the transmission medium, such as an optical fiber. If we launch power into a spatial beam form that is a combination of modes of different phase velocities in the fiber, then the beam form will change as it propagates down the fiber. That change in beam shape is not itself a problem for this add-drop device – it will still align itself to that beam shape. If there is also dispersion in the fiber so that, with sufficiently different frequencies, the arriving beam shape is substantially different for different frequencies, then we need to use different SRADMs for different frequencies, separating those frequencies before the SRADMs (e.g., by putting wavelength splitters in each of the waveguides W11 – W14 to separate to different SRADMs and corresponding wavelength combiners in the waveguides WO1 – WO4 to combine the outputs from different SRADMs). If the received beam shape does not vary substantially with wavelength, we can use one SRADM, and we could instead put wavelength splitters before the receivers, and wavelength combiners after the transmitters, if we wanted to have separate wavelength channels.

Another important question with such a scheme is how fast such a self-aligning scheme needs to adapt to changing conditions. For fibers in which there is negligible mode coupling during propagation, such as short multimode fibers using their lowest modes or short few-mode fibers, the mode shapes to be separated will not be changing. If, however, we use long fibers, the speed at which we may be able to run the necessary feedback loops may become an issue. For example, recent work analyzing 50 km of few-mode fiber shows fluctuations in the range of 10's of kHz and optimum adaptation rates for MIMO equalization in the 100's of kHz. Though it may be possible to make quite fast analog electronic feedback loops, we note that to self-align to a beam with N segments requires N successive feedback loops to settle in the architecture shown here. (An alternate architecture given in [26] with a binary tree format of Mach-Zehnder interferometers requires only $\log_2 N$ successive feedback loops for optimization of one beam.) Hence, whether such a scheme is viable for use with long fibers with mode coupling is an open question on these grounds. Use with such long fibers may also be problematic because different group delays on different propagating modes may mix different symbols in a bit stream.

In this discussion, we have used only 4 waveguides and 4-element input and output couplers as an illustration. The SRADM concept is simply extended to larger numbers of elements. How many elements we need in a given situation depends on the complexity of the beams we are working with. The issue of the necessary complexity for an optical component to perform a specific function has been discussed generally in [33]. If our device is to separate one mode from M possibilities (and hence also to pass $M - 1$ other modes through the device), we need to have at least M beam coupling elements (e.g., grating couplers) and M waveguides with associated adjustable beam splitters and phase shifters (e.g., MZIs).

3. Use in multiple channel systems

3.1 Setup in the presence of multiple active spatial modes

So far, we have discussed setting up the SRADM device when only the add-drop mode of interest is present at the input. In a real communications system, it would be very desirable be able to set up the SRADM up when possibly all of the spatial modes were in use. This could be achieved here by a simple additional coding on the add-drop channel [26]. For example, at the original source of the various spatial channels, we could impose specific small low-frequency power modulation on the signals, at a different frequency for each channel. Then, to lock on to a specific channel, we would look for signals coming out of the detectors D1 – D3 with that frequency component, which could be achieved with a narrow band electrical

filter or a lock-in amplifier technique [34]. Then power in other orthogonal modes would be ignored by these detectors, allowing the feedback loops to operate on the desired add-drop mode even in the presence of other modes. Note that such a lock-in technique only has to operate at a presumed relatively slow speed for the self-alignment process, not the signal bandwidth, and it only needs to monitor the magnitude of the lock-in signal, not the phase.

Note that this approach allows the overall multimode communications channel to have arbitrary scattering between different orthogonal spatial modes (e.g., in an optical fiber) as long as that scattering itself is linear and loss-less (or equal for all modes) and can therefore be represented by a unitary operator (at least within a multiplying constant). Such unitary scattering (at least within a multiplying constant) retains orthogonality between channels even if each channel is now represented by a different spatial function, and this device could therefore still lock on to the desired channel and pass the others even if the spatial modes themselves had become mixed in such a unitary fashion. (See Ref [28]. for a discussion of how to handle non-unitary linear scattering between modes.) It is important, however, to take note of any different propagation delays in different modes as these would cause problems with mixing of symbols from different bit periods.

3.2 Bypass and hitless operation

We could extend the SRADM device to add a “bypass” function. In such a bypass case, the signal that would normally be dumped into the receiver is instead routed round the receiver and transmitter, just like routing a train round a station. This could be accomplished by adding appropriate routing switches (such as additional Mach-Zehnder interferometers) in front of the receiver and just after the transmitter, together with a bypass waveguide, as shown in Fig. 3.



Fig. 3. Bypass switching. Additional waveguide switches, here implemented using MZIs ML1 and MR1, can be used to connect waveguide WA1(I) directly to WA1(O), bypassing the receiver R1 and transmitter T1.

We could use this bypass setting to set up the SRADM without interrupting any of the channels going through the device (i.e., “hitless” operation [35]). (Note that for truly hitless operation, it may be necessary to use “endless” phase shifters [36] in the device so there are not discontinuous jumps as the end of the range is reached for a given phase shifter.) In this bypass setting, we can use the detectors to set up the internal state of the device without changing the fact that all modes are being transmitted straight through the device. (Note explicitly that, in this bypass mode, as long as the MZIs MI1 – MI4 and MO1 – MO4 are always set with the same “reflectivities” and the opposite phase shifts, respectively, as discussed above, the SRADM makes no change in the transmission of any particular mode.) Then, once we have set the device, using the detectors and feedback loops, so that the add-drop mode now of interest to us is passing through the top waveguide, we can switch in the receiver and transmitter instead of the bypass waveguide. Hence, the device can be optimized in preparation for adding and dropping any specific mode while passing all the modes unchanged through the device, making this setup totally “hitless”. Once we have set the device in preparation for a given add-drop mode, we can switch back the bypass switches, thereby connecting the add-drop receiver R1 and transmitter T1 back into the system, with no change in the transmission of the other modes of the system.

3.3 Simplified device

If we do not mind the spatial channels being of different form in every part of the network, we could possibly use a simplified form of the SRADM, as shown in Fig. 4. Such an approach

might be viable if the modes in use in the network all have similar delay (e.g., if they are part of the same mode group.) This device is simply a self-aligning mode coupler [26] implemented with mostly-transparent detectors and output waveguides on the far side of those detectors and, optionally, with additional dummy MZIs for equal loss. Here we will presume that there is some way for the detectors to identify the channel of interest, such as a low-frequency modulation of that channel to allow the detectors to lock-on to that channel only, as discussed above.

This device will certainly change the mode forms of the different channels as they pass through but, as long as any scattering between modes during propagation in the network is unitary (within a multiplying constant), the various spatial channels remain orthogonal and can still be separated by other similar devices down-stream, again on the presumption we have some labeling, such as a different low-frequency modulation for each channel, that can be recognized by the detectors. In such a scheme, it is only important that the detectors can recognize the channel of interest and that the channels remain orthogonal, as will be the case if the scattering in the system is unitary (lossless) or unitary within a single loss factor.




Fig. 4. Simplified device using only a self-aligned mode coupler configuration [26] with added “through” waveguides and mostly-transparent detectors. We could also add dummy MZIs as in the left side of Fig. 2 so all paths have equal numbers of MZIs for loss and path length equality.

3.4 Universal device

We could extend the concepts here to make a SRADM that can add and drop any channel or combination of channels, all in a hitless fashion, as shown in Fig. 5. In this device, we have added in receivers (R1 – R4), transmitters (T1 – T4), and pairs of bypass switches (ML1 & MR1 – ML4 & MR4) in each path. We have also added additional diagonal “rows” of MZIs and mostly-transparent detectors to allow the simultaneous separation of each of the different orthogonal modes; this configuration corresponds to the self-aligning coupler [26] when configured to align to multiple beams at once. The MZIs in the first row, MI11 – MI14, are set using the signals in detectors D11 – D14, exactly as before, to put one of the desired modes into the R1 – T1 channel. Then, similarly, MI21 – MI23 are set using the signals in detectors D21 – D22 to put the second mode into the R2 – T2 channel. Finally, in this example, MI31 – MI32 are set using the signal from detector D31 to put the third mode into the R3 – T3 channel, with the fourth mode being automatically sent to the R4 – T4 channel as a result.

The MZIs MO11 – MO32 on the right hand side are set with the same reflectivities but opposite phases to their similarly numbered counterparts on the left. The modes as separated in the center can either be switched for add-drop or for bypass by using the bypass switch pairs (ML1 & MR1 – ML4 & MR4) as desired. If all the channels are set for bypass, the entire SRADM device can be reconfigured in a hitless fashion.




Fig. 5. Universal SRADM device configuration that allows add-drop of any number or combination of the channels as well as hitless switching operation. Receivers (R1 – R4), transmitters (T1 – T4), and pairs of bypass switches (ML1 & MR1 – ML4 & MR4) have been inserted in each path, and additional diagonal rows of detectors and MZIs have been added to allow simultaneous separation of multiple modes. The devices on the right are set to the same “reflectivities” but opposite phase delays compared to their corresponding devices on the left.

To use the different rows of detectors to set the different rows of MZIs while all the beams are illuminated simultaneously in the device, we need to use a technique that enables us to distinguish the different mode signals in the detectors, such as using a different low-frequency modulation on each mode [26] as mentioned above. Detectors D11 – D13 can be set to look for the low-frequency modulation of the first mode, D21 – D22 to look for the low-frequency modulation of the second mode, and D31 for the low-frequency modulation of the third mode.

4. Conclusions

We have shown here how to make an add-drop multiplexer for arbitrary spatial modes. A SRADM device of this type could be implemented using any of several different approaches to planar optical circuits, including thermo-optic silicon-based technologies [37], for example, as long as appropriate mostly-transparent detectors are also available (see [26,27] for discussions of approaches). The concept is self-aligning based only on local feedback loops where signals from mostly-transparent detectors are used to set controllable beam-splitters and phase shifters. Note that there is no global multiparameter optimization required in this approach; the device steps progressively through multiple local feedback operations based on minimizing power in detectors. By using some simple coding to identify each spatial communication channel of interest, such as a low-frequency power modulation at a different frequency for each different channel, the SRADM can optimize its alignment to one spatial channel even in the presence of power in the other channels. The alignment optimization can be left running continuously while the SRADM is in use so that it can even compensate for changes in the spatial modes. By choosing equal path lengths for all the beam paths through the system, the SRADM itself can be substantially independent of wavelength, with one setting working for many different wavelengths. The device can operate in a hitless mode that allows the device to be reconfigured without disturbing the transmission of any of the channels.

This approach shows that we may be able to make use of multiple spatial modes flexibly in complex transmission systems, and may open the way for broader application of spatial modes in communications.

Acknowledgments

This project was supported by funds from Duke University under an award from DARPA InPho program, and by Multidisciplinary University Research Initiative grants (Air Force Office of Scientific Research, FA9550-10-1-0264 and FA9550-09-0704).