All-optical crossbar switch using wavelength division multiplexing and vertical-cavity surface-emitting lasers

Brian Webb and Ahmed Louri

A design for an all-optical crossbar network utilizing wavelength-tunable vertical-cavity surface-emitting laser (VCSEL) technology and a combination of free-space optics and compact optical waveguides is presented. Polymer waveguides route the optical signals from a spatially distributed array of processors to a central free-space optical crossbar, producing a passive, all-optical, fully connected crossbar network directly from processor to processor. The analyzed network could, relatively inexpensively, connect local clusters of tightly integrated processors. In addition, it is also believed that such a network could be extended, with wavelength reuse, to connect much larger numbers of processors in a multicluster network. © 1999 Optical Society of America

OCIS codes: 200.4650, 200.2610, 090.1970, 140.3600, 250.7260, 060.4230.

1. Introduction

Although great advances are being made to increase the speed and extend the range of high-speed electronic interconnects, it is becoming increasingly more difficult to keep pace with the bandwidth and connectivity requirements of modern microprocessors and multiprocessor systems. Optics offers many advantages over electronics that add great potential for high-bandwidth multiprocessor interconnects, such as large-bandwidth capability, low-power requirements, relative immunity to electromagnetic interference, and three-dimensional free-space propagation. Unfortunately, optical interconnect systems have traditionally been relatively complex and difficult to implement in a cost-effective manner. Recent advances in micro-optical components and interconnect technologies provide some advantages that could be used to develop high-bandwidth multiprocessor interconnects with a relatively small number of components, with technology that matches well with traditional printed-circuit- (PC-) board-based systems. It is believed that these advances in optoelectronic technology could provide some advantages that make optics a more viable replacement for electronic interconnects.

One technology that is becoming increasingly more common in optical communications systems is the vertical-cavity surface-emitting laser (VCSEL).[1–3] Recent advances in VCSEL technology have added the ability to produce VCSEL's with wavelengths that can be externally tuned to a wide range of frequencies.[4–8] Wavelength-tunable VCSEL's allow for multiple wavelengths to be used for wavelength division multiplexing (WDM),[9–12] which can reduce the number of VCSEL's required per processor and can reduce the complexity of the interconnection network by multiplexing multiple signals over a single transmission line. Using tunable VCSEL's with WDM facilitates building highly scalable, compact, dense interconnection networks that use less optical components than do traditional optical networks with fixed-wavelength technologies.

In this paper we propose and analyze a crossbar interconnect implemented with tunable VCSEL's, using WDM. A WDM crossbar can be constructed with a single tunable VCSEL and a single fixed-frequency receiver per processor. The signals from all tunable VCSEL's are routed to a passive WDM optical crossbar that uses a free-space, single-grating crossbar–demultiplexer to route the signals to the correct destination processor. With optical WDM the number of components required for creating such a crossbar network scales as a factor of $O(N)$, where N is the

The authors are with the Department of Electrical and Computer Engineering, the University of Arizona, Tucson, Arizona 85721. A. Louri's e-mail address is louri@ece.arizona.edu.

Received 14 December 1998; revised manuscript received 11 May 1999.

0003-6935/99/296176-08$15.00/0

number of processors, whereas traditional electronic crossbars require $O(N^2)$ switches and wires to implement a crossbar, which tends to limit its scalability.

One problem associated with using VCSEL-based free-space optical interconnects is that a free-space optical system should be compact, owing to the feature sizes of the optical components and alignment constraints of the optics. To solve this problem, we also propose the use of polymer waveguides[13–16] to route the optical signals from the physically separated discrete processors to a centralized free-space optical crossbar. Polymer waveguides can be fabricated directly on a PC board, thus allowing for the processors to be distributed around the PC board, closely matching the traditional design of PC-board-based computer systems. Polymer waveguides can also be used in optical backplanes to route signals among multiple PC boards to construct larger-scale systems. An advantage to using polymer waveguides is that the optics can be tightly integrated with the processors, and the processors can be distributed over relatively large distances, thus creating a high-bandwidth, all-optical network between multiple discrete processors and possibly multiple PC boards. In addition, utilizing WDM along the waveguides greatly reduces the number of waveguides required.

The proposed WDM VCSEL-based crossbar features the following characteristics: (1) It fully exploits the bandwidth and cost advantages of WDM and the inherent benefits of three-dimensional free-space parallel optics; (2) it supports gigahertz bandwidths and channel switching times in the megahertz range; (3) it is relatively simple, reducing design complexity and alignment difficulties and reducing the number of limits (diffraction limits, latencies, and power losses) imposed by the optical components; (4) there is no optical power loss, owing to optical fanout; and (5) integrated polymer waveguides are used to route spatially separated processors to a centralized optical crossbar, implementing an all-optical connection directly from processor to processor, which makes the overall crossbar design and implementation compatible with current processor/PC-board-based technologies, which reduces the difficulty in migrating to optical interconnect technologies. It is believed that this optical interconnect could be used to construct inexpensive optical interconnects for medium-scale parallel processors and that it could be extended through wavelength reuse[17,18] to support much larger highly parallel architectures.

2. Optical Implementation of a Wavelength Division Multiplexing Optical Crossbar

An overview of the proposed optical crossbar can be seen in Fig. 1. Each processor contains a single tunable VCSEL that is tightly integrated with the processor. Optical waveguides are used to route the optical signals from each processor to a central free-space optical crossbar. The all-optical crossbar uses WDM, so only a single transmit–receive waveguide pair is required per processor. The transmit waveguides from each processor are coupled into a single waveguide by use of a passive optical combiner, and the combined optical signal is then routed to a grating-based free-space optical crossbar–demultiplexer that routes the optical signals to the appropriate output waveguide. These waveguides are then routed back to an integrated optical receiver.

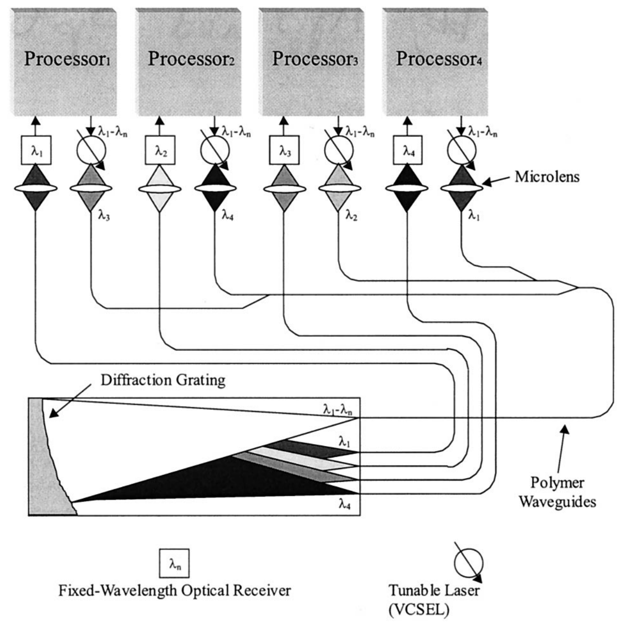

Fig. 1. Conceptual overview of a free-space optical crossbar with WDM for interprocessor interconnects. Tunable VCSEL lasers are used to select the destination processor, and each processor contains a fixed-wavelength optical receiver. A concave diffraction grating is used to diffract individual wavelengths to the appropriate processor.

A. Wavelength-Tunable Vertical-Cavity Surface-Emitting Lasers

The first stage of the optical crossbar is the wavelength-tunable VCSEL's. Promising research into wavelength-tunable VCSEL's is currently being conducted by researchers at Stanford University[4–8] as well as others.[19,20] The Stanford researchers are producing VCSEL's that include a deformable membrane as the top mirror of the laser cavity. This membrane can be adjusted continuously with electrostatic charge, which adjusts the dimensions of the laser cavity. Adjusting the length of the laser cavity produces a corresponding lengthening of the optical wavelength output from the laser. These tunable VCSEL's are designed to have a base wavelength of ~960 nm and a continuous-wavelength tuning range that has recently reached 31.6 nm.[8] The spectral width of these VCSEL's is less than 1 nm, so a spectral resolution of 1 nm is possible. A WDM crossbar using these VCSEL's and a 1-nm channel spacing could support approximately 32 channels, or 16 channels with a 2-nm channel spacing. The power output ranges from 0.56 to 1.6 mW at a 10-mA bias current.[8]

These VCSEL's can be tuned in the megahertz frequency range[4]; so tuning speeds of the order of 1 μs are possible. This tuning speed is relatively fast for a micromechanical tunable optical device, but it may become a limitation for high-bandwidth communications in parallel computers, where tuning latency is more of an issue. Tunable VCSEL research is still in the early stages, though; so it would seem reasonable to expect the tuning speeds to improve. If higher-speed tuning is required, another alternative is to use multiwavelength VCSEL arrays,[21,22] in which the single tunable VCSEL is replaced with multiple fixed-frequency VCSEL's that each transmit at a slightly different wavelength. This alternative increases the number of VCSEL's required for implementing the crossbar, but the optical signal from all VCSEL's can be combined into a single waveguide; so the rest of the crossbar hardware remains the same.

B. Polymer Waveguides

It is possible to implement an all-optical crossbar with free-space optics.[23] A free-space optical system has some advantages over a waveguide-based optical system. Free-space optical systems can better utilize the three-dimensional nature of optics and can implement compact and complex interconnection patterns with a minimum of interconnection hardware. One problem associated with free-space optical interconnects, though, is that a free-space optical system should be compact, owing to the feature sizes of the optical components and the optical alignment constraints, but the components being connected (discrete processors, and the like) are not necessarily compact and cannot necessarily be arranged in the configuration required for the free-space optical system. Therefore it seems advantageous to find a method to route the optical signals from the physically separated discrete processors to a centralized free-space optical crossbar. Polymer waveguides[13–16] seem to provide a good means for achieving this goal. Polymer waveguides can be fabricated directly on a PC board, allowing for the processors to be spaced out around the PC board. The waveguides route the optical signals into and out of the optical crossbar and facilitate building an all-optical interconnect system.

Using polymer waveguides to route the optical signals to the optical crossbar allows for the optical transmitters and receivers to be integrated directly on the processor. This has the advantage of reducing the parasitic capacitance, allowing for higher potential bandwidths, lower potential latencies, and reduced power requirements. There are several methods for integrating VCSEL's and optical receivers with standard complementary metal-oxide silicon circuitry[24,25] including substrate removal,[25] coplanar flip-chip bonding[24] (for bottom-emitting VCSEL's), and various top-contact bonding methods.[24] The processor integrated VCSEL's can be coupled with PC board integrated polymer waveguides with standard butt-coupling techniques or with a microlens coupler (as depicted in Fig. 1).

Polymer waveguides can be constructed in various sizes and pitches from a few micrometers to tens of micrometers or more. Multimode polymer waveguides have been constructed with a 50 μm × 50 μm core and a 100-μm pitch[15] for use with PC board and optical backplane applications, with large numbers of these waveguides fitting into a compact space. The size and spacing of these polymer waveguides match well with the feature size and spacing of VCSEL's and optical receivers. The 100-μm pitch is also a reasonable spacing for integrating with the free-space WDM demultiplexer, although a smaller waveguide size and pitch could be supported, which would make for a more compact crossbar.

C. Optical Combiner

The individual polymer waveguides from each processor route the optical signals from the processors to a centralized optical combiner that couples the signals from the individual waveguides into a single waveguide. This produces a single multiwavelength optical signal that is routed to the free-space optical crossbar–demultiplexer. It is not difficult to construct an optical combiner network with polymer waveguides. 2 × 1 Y couplers can easily be constructed in polymer waveguides,[15,16,26] and these Y couplers can be combined to construct a binary tree optical combiner network. A binary tree combiner successively combines the signals from each input source into a single optical waveguide. The 50 μm × 50 μm core size and 100-μm pitch of the polymer waveguides make it possible to construct relatively large optical combiner networks in a small space directly on the PC board.

D. Free-Space Optical Crossbar–Demultiplexer

The combined optical signal on the single waveguide output from the optical combiner contains multiple individual wavelength channels that are destined for different processors. Each processor is assigned a fixed wavelength with which to receive, and other processors can transmit to that processor by transmitting on the wavelength assigned to the receiving processor. For example (see Fig. 1), for processor 1 to transmit to processor 3, processor 1 would simply transmit on the wavelength assigned to processor 3 (e.g., λ_3). This WDM scheme requires tunable transmitters, fixed receivers, and some way to route each signal to the correct destination processor.

Two common designs that are employed for demultiplexing WDM networks are (i) broadcast and select and (ii) wavelength routing.[27] The former is implemented by broadcasting of the optical signal to all processors, with the receiving processor somehow selecting only the signal transmitted on its own wavelength. A wavelength-routed network, however, uses passive or active optics to route the individual wavelengths to the appropriate processors. The crossbar described in this paper uses wavelength routing to demultiplex the optical signals.

A wavelength-routed crossbar network can be constructed with a standard diffraction-grating-based demultiplexer, in which the combined input signal is

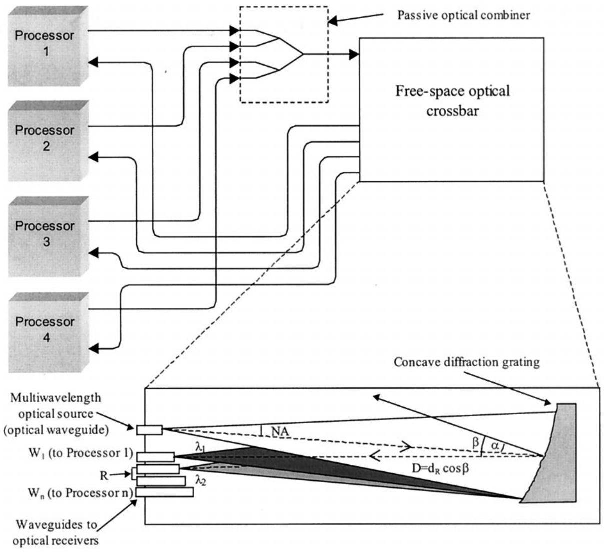

Fig. 2. Proposed compact optical crossbar consisting of polymer waveguides directly coupled to processor-mounted VCSEL's, a polymer waveguide-based optical combiner, and a free-space optical crossbar–demultiplexer based on a concave sawtooth diffraction grating.

separated into the individual wavelength channels by means of a diffraction grating. There are two popular grating types used in wavelength-multiplexed systems, sawtooth gratings and holographic gratings.[28] Sawtooth gratings can approach 100% diffraction efficiency to the first diffraction order and can be produced fairly readily by means of etching or by photolithographic processes. Holographic gratings can be produced with a photographic process and can also approach 100% diffraction efficiency if the Bragg phase-matching condition is satisfied. In this paper we analyze the use of a concave sawtooth grating to perform both the function of wavelength separation and that of focusing the beam from the input waveguide to the various output waveguides.

An overview of a concave-diffraction-grating-based demultiplexer can be seen in the inset in Fig. 2. The input to the demultiplexer is a single optical waveguide containing the combined optical signals. The output is multiple optical waveguides that route the separate wavelength channels to the individual processors (one per processor). The diffraction grating must be designed and aligned such that, when the optical source is tuned to the longest wavelength λ_1, the first diffraction order will fall on the first output waveguide in the output waveguide array W_1. When the optical source is tuned to a shorter wavelength λ_n, the beam should fall on the nth output waveguide W_n.

For a concave diffraction grating based on the Roland circle,[28] the focus condition is defined by the equation

$$\Lambda(\sin \alpha + \sin \beta) = n\lambda, \tag{1}$$

where Λ is the grating period at the center of the grating, β is the angle of incidence of the input beam with respect to the normal to the center of the grating, α is the angle of the diffracted beam with respect to the normal, and n is the diffraction order (1 in this case).

The angular dispersion of the grating can be found by differentiation of Eq. (1) with respect to λ:

$$\frac{\Delta\beta}{\Delta\lambda} = \frac{n}{\Lambda \cos \beta}. \tag{2}$$

For the first diffraction order ($n = 1$), and for small variations in the wavelength λ, this equation defines a linear relationship between $\Delta\lambda$ and $\Delta\beta$. If we term R as the linear offset in the output plane between beams of two different wavelengths and D as the distance from the optical grating to the output plane, for small $\Delta\beta$ there is a linear relationship between the change in the wavelength of the beam $\Delta\beta$ and the offset R:

$$\Delta\beta = \frac{\Delta\lambda}{\Lambda \cos \beta} = \tan^{-1}\left(\frac{R}{D}\right) \cong \frac{R}{D}. \tag{3}$$

To resolve the individual channels at the output waveguide array, the focused spot size at the output waveguides must be at least as small as the core of the output waveguides, and the angular dispersion between adjacent channels must be equal to the pitch between neighboring waveguides. From Eq. (3),

$$D = \frac{R}{\Delta\beta} = \frac{R\Lambda \cos \beta}{\Delta\lambda}. \tag{4}$$

To produce a compact system, we would prefer to resolve the individual channels in the shortest possible distance. This requires reducing the grating period to as small as possible. To reduce aberrations, the grating period should be no smaller than a few wavelengths of the optical source.[29] Since the VCSEL's have a base wavelength of approximately 1 μm, we will assume a grating period of 2 μm and a channel spacing of 2 nm. If we choose an incident angle from the normal to the center of the grating of $\alpha = 13° = 0.242$ rad, we get a base diffraction angle β from the normal to the center of the grating of [from Eq. (1)]

$$\beta = \sin^{-1}\left(\frac{\lambda}{\Lambda} - \sin \alpha\right) = 0.243 \text{ rad} \tag{5}$$

and angular dispersion $\Delta\beta$ for the first diffraction order of [from Eq. (2)]

$$\Delta\beta = \frac{\Delta\lambda}{\Lambda \cos \beta} = 1.0 \text{ mrad}. \tag{6}$$

Therefore, if we assume a waveguide pitch of 100 μm, using Eq. (4) we can calculate the distance from the center of the optical grating to the output plane required for resolving a 2-nm change in wavelength:

$$D = \frac{R}{\Delta\beta} = 10 \text{ cm}. \tag{7}$$

The diameter of the Rowland circle d_R can then be calculated as

$$d_R = \frac{D}{\cos \beta} = 10.3 \text{ cm}, \tag{8}$$

and the grating should be constructed with a radius of curvature of 10.3 cm.

Assuming a multimoded numerical aperture of 0.29 for the input polymer waveguide, a grating of approximately 6 cm × 6 cm would be required for capturing and refocusing the expanded beam. If a waveguide with a smaller numerical aperture is used, the size of the grating could be reduced accordingly.

E. Diffracted Spot Size and Cross-Talk Estimation

The primary potential source of cross talk in this free-space, grating-based optical crossbar is overlap of neighboring spots in the spread spectral diffraction pattern of the diffraction grating. The spot size must be smaller than the pitch of the output waveguides, or there will be excessive cross talk between neighboring channels. If the outputs are waveguides, the angle of the incident beam should also closely match the numerical aperture of the output waveguide. In this case the input and the output waveguides are spatially close; so, if the input and the output waveguides are similar, the angle of incidence of the beam should closely match the numerical aperture of the output waveguide.

The spot size of the diffracted beams will be governed first by the diffraction limit of the optics and second by aberrations and imperfections in the grating. If we assume the optical waveguide to be approximately an $f/2$ system, the diffraction-limited spot diameter is approximately 4.7 μm. This is much smaller than the 100-μm spacing assumed between the output optical waveguides. Other factors that contribute to spreading the beam at the output plane include spherical aberrations caused by nonoptimal spacing of the grooves and astigmatism caused by the shape of the grating.[28] Imperfections in the grating grooves, alignment error, and other imperfections can also cause spreading of the spot at the output plane.

Similar concave grating demultiplexers have been implemented. For example, a similar concave grating demultiplexer has been implemented to demultiplex 0.4-nm separated channels spaced around a central wavelength of approximately 1.5 μm.[30] This demultiplexer successfully separated signals from a single-mode ($f/5$) optical fiber to multiple single-mode output fibers and achieved a spot size of approximately 20 μm, which is near the diffraction limit for an $f/5$ system. In this demultiplexer, the output waveguides were separated by 42 μm and the cross-talk isolation between adjacent channels was measured at greater than 25 dB. If this is assumed to be typical of such a grating demultiplexer, it would seem reasonable to assume a cross-talk isolation within the grating demultiplexer in this crossbar of greater than 25 dB.

3. Power Analysis and Bit-Error-Rate Estimation

Calculation of a power budget and the signal-to-noise ratio at the receiver is important for confirming the realizability and scalability of an optical interconnect implementation. The signal-to-noise ratio at the receiver gives an indication of the expected bit-error rate (BER) of the digital data stream. For optical communications networks it is acceptable to have a BER of 10^{-9} or greater, but for interprocessor optical interconnect networks it is standard to require a BER of 10^{-15} or greater.

A. Required Optical Power

The BER of an optical system can be calculated as follows:

$$\text{BER} = \frac{1}{\sqrt{2\pi}} \frac{\exp(-Q^2/2)}{Q}, \tag{9}$$

where Q is twice the signal-to-noise ratio required at the receiver to provide the specified BER. For example, a BER of 10^{-9} corresponds to $Q = 6$, and a BER or 10^{-15} corresponds to $Q = 7.94$.

If we assume that a BER of 10^{-15} is required, we can calculate an estimation of the minimum power required at the receiver to provide the given BER as follows[23]:

$$\bar{P} = \frac{2\pi hc}{q\eta\lambda} \left(\frac{4kT\Gamma I_3}{g_m} \right)^{1/2} Q C_T B^{3/2}, \tag{10}$$

where h is Planck's constant, c is the speed of light, q is the electron charge, η is the detector quantum efficiency, k is Boltzmann's constant, T is the absolute ambient temperature, Γ is the field-effect transistor (FET) channel noise factor, I_3 is a weighting function that is dependent on the input optical pulse shape, g_m is the FET transconductance, C_T is the total capacitance in the receiver circuit, and B is the signal bandwidth.

If we assume a GaAs metal–semiconductor FET receiver with a quantum efficiency η of 80%, a FET channel noise factor Γ of 0.7, an optical pulse weighting function I_3 of 0.0868, a FET transconductance g_m of 30 mS, a total capacitance C_T of 0.75 pF, and a data rate of 2 Gbit/s,[31] we can achieve a BER of 10^{-15} with an optical power at the receiver of 1 μW or −30 dBm.

B. Optical Losses and Power Budget

To determine whether the required power level will be present at the optical receiver, it is necessary to estimate the optical power emitted by the VCSEL's and estimate the losses incurred throughout the optical components of the system. At a 10-mA bias current, wavelength-tunable VCSEL's output between 0.56 and 1.6 mW of optical power. For this analysis we will assume a value at the high end of this range, because the tunable VCSEL's are still in the early research stages, and it is likely that the

Table 1. Losses for Each Component of the Optical Crossbar

Loss Mechanism	Loss (dB)
VCSEL–waveguide coupling (L_{vc})	-1
Waveguide (L_w)	-6
Y coupler (L_Y)	$-3\log_2(n)$
Demultiplexing (L_d)	-9
Receiver coupling (L_{rc})	-0.5
Total	$-16.5 - 3\log_2(n)$

optical power output will increase as the technology improves; so we will assume a starting optical power output of 1.6 mW or 2.04 dBm.

The total optical loss in the system is the sum total of the losses (in decibels) of all optical components that a beam must pass through from the transmitter to the receiver. Optical losses are incurred in the following components:

- *VCSEL-waveguide coupling* (L_{vc}). There will be losses incurred during coupling of the beam emitted from the VCSEL's into the optical waveguides. The insertion loss for a commercially available fiber coupler is taken as -1 dB, so we will assume a -1-dB loss for VCSEL–waveguide coupling.
- *Waveguide* (L_w). Polymer waveguides have a much higher absorption loss than do silicon waveguides, so the absorption losses within the polymer waveguides must be accounted for. This will include the losses incurred both while the beam is routed from each processor to the optical crossbar and while the beams are routed from the optical crossbar back to the receivers. Polymer waveguides are being constructed with losses of less than -0.1 dB/cm.[15] It has been observed for similar waveguides that the losses for 90° bends with radii of curvature of greater than approximately 1 mm are approximately -0.1 dB.[32] We will assume a total length of 50 cm of polymer waveguide from each transmitting VCSEL to the receivers, and we will assume that each waveguide will contain no more than approximately ten 90° bends. Although for large crossbar networks this length will vary with the size of the network; this should be a good maximum estimation for a system containing a reasonable number of processors, so the total loss within the polymer waveguides is assumed to be -6 dB.
- *Y coupler* (L_Y). We will assume a loss of -3 dB/Y coupler, which gives a total optical loss in the optical combiner network of -3 dB $\log_2(n)$.
- *Demultiplexing* (L_d). Losses will be incurred from various sources within the grating demultiplexer. Losses will be incurred owing to imperfections in the grating, diffraction efficiency losses from the grating, and coupling losses that are due to scattering, waveguide coupling mismatches, and Fresnel losses. A similar grating demultiplexer[30] was observed to have losses ranging from -6 to -9 dB for 40 optical channels; so we will assume a loss within the free-space demultiplexer of -9 dB.

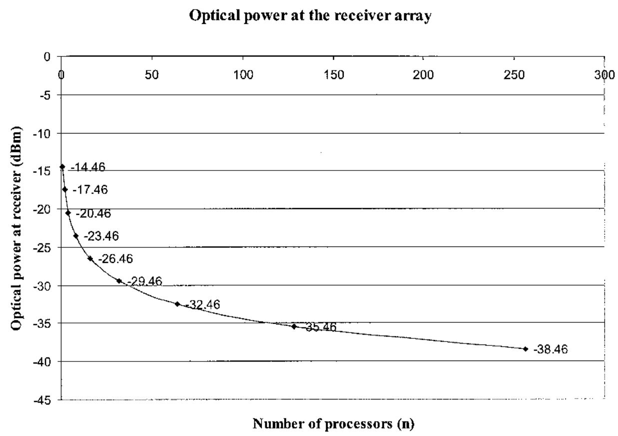

Fig. 3. Optical power at each optical receiver for varying numbers of processors (channels). It can be seen that, from a power budget perspective, the -30-dBm optical power required by the receivers is sufficient to support 32 processors, although improvements in the VCSEL efficiencies would likely increase the maximum number of processors.

- *Receiver coupling* (L_{rc}). Some losses will be incurred when the beam is coupled from the polymer waveguides onto the optical receiver. We will assume a loss of -0.5 dB.

Assuming an n input $\times$ n output optical crossbar, the individual losses are estimated in Table 1. The total transmission loss for the optical crossbar network is the total of all the losses:

$$L_{\text{total}} = L_{vc} + L_w + L_Y + L_d + L_{rc} = -16.5 \text{ dB} + -3 \text{ dB} \log_2(n). \quad (11)$$

It should be noted that the only size-dependent parameter of the loss equation is the losses in the tree combiner network. All other losses are fixed and will not vary with the size of the network (except for the length of the polymer waveguide, which is discussed above).

If we assume a VCSEL power of 2.04 dBm, we can plot the resulting optical power at the receivers for varying crossbar sizes (Fig. 3). It can be seen from Fig. 3 that 32 processors (channels) can be supported with a BER of less than 10^{-15}. It is expected that further advances in tunable VCSEL technologies should increase the power efficiencies of these VCSEL's, which would allow for larger crossbar sizes.

4. Integration of the Free-Space Optical Crossbar into Larger Systems

One of the limiting factors of WDM is the limit of the number of wavelength channels that can be supported by the optical components in the system. Several factors limit the number of channels supported, including the tunability range of the VCSEL's, the thermal stability of the VCSEL's, the spectral width of the VCSEL beam, the resolving power of the grating demultiplexer, grating diffraction limits, and the like. If a WDM interconnect

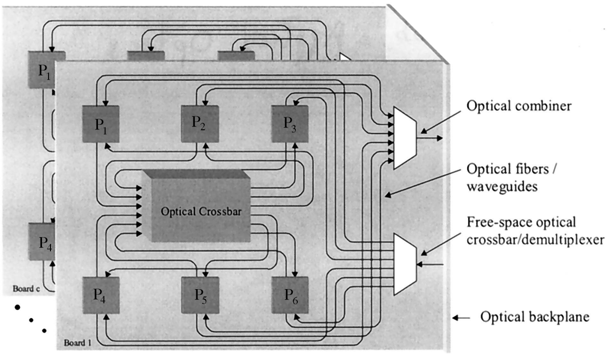

Fig. 4. Proposed implementation of a hierarchical network architecture called the SOCN. A SOCN system consists of a network of multiprocessor clusters linked by means of free-space and fiber-based WDM optical crossbars. Free-space optical interconnects are used for intraboard interconnects (within the cluster), and fiber interconnects are used for interboard interconnects (between clusters). Polymer waveguides are used to route the optical signals on the PC boards.

relies solely on a single set of wavelength channels, then ultimately compromises will have to be made when a large number of processors are supported. One way to extend a WDM optical interconnect beyond the limits of the number of channels supported is by wavelength reuse.[17,18] As the name implies, wavelength reuse optically isolates portions of the interconnect network to allow for the use of the same set of wavelengths simultaneously in multiple parts of the system for different purposes.

One method of utilizing wavelength reuse to extend a WDM interconnected system is to connect various portions of the system by means of optically separated networks. In the context of this free-space optical crossbar, wavelength reuse could be utilized by the addition of more VCSEL's and optical receivers to create multiple network connections on each processor. Since VCSEL's and optical receivers are compact, a reasonably sized array of VCSEL's and receivers can be fitted onto a modern microprocessor. Each VCSEL–receiver pair could be thought of as an independent network interface, connecting the processor to different parts of the system.

In Fig. 4 we propose an implementation of a hierarchical network architecture called the scalable optical crossbar network (SOCN), consisting of clusters of processors contained on individual PC boards. Each processor on a cluster contains a VCSEL–receiver pair that is connected to a local (intracluster) free-space optical crossbar by means of polymer waveguides. Each processor in a SOCN cluster also contains another set of multiple VCSEL–receiver pairs that connect the processors on one cluster to the processors on other clusters, creating a set of intercluster interconnects. For any given intercluster connection, the VCSEL corresponding to that connection is connected to an intercluster optical combiner, which could be constructed similar to the optical combiner used for the intracluster free-space optical crossbar, by means of polymer waveguides. The output of the optical combiner is coupled with a standard optical fiber, which is routed to the remote cluster, possibly through an optical backplane configuration. The fiber is then coupled into a waveguide on the remote cluster that is routed to an optical demultiplexer. This optical demultiplexer could also be constructed similar to the intercluster crossbar–demultiplexer. The demultiplexed signals are then routed to the appropriate remote processor by use of polymer waveguides. This, in effect, extends the local free-space optical crossbar, with wavelength reuse, over optical fibers to remote clusters, potentially creating much larger systems.

If there are m clusters in a system, each containing n processors, and if each cluster contains an intercluster link to every other cluster, a full $N = n \times m$ processor crossbar network could be created. To construct such a crossbar, each processor would have to contain m VCSEL's and m optical receivers. Each cluster would contain one intracluster crossbar network and $m - 1$ intercluster crossbar connections. If a fully connected network is not required, it should also be possible to extend the system further by connection of the clusters by means of an intercluster network other than a fully connected crossbar, thereby trading full connectedness for a larger system size. Details of the proposed SOCN architecture will be the subject of a separate study.

5. Conclusions

We have presented a design for an all-optical, wavelength division multiplexed (WDM) crossbar interconnect that uses wavelength-tunable VCSEL's and is integrated with polymer waveguides. On the basis of already achieved results from tunable VCSEL research, it should be possible to construct a fully connected optical crossbar network that will support at least 16 channels, assuming a 2-nm channel spacing, or 32 channels with a 1-nm channel spacing. As the technology improves, the number of channels supported should increase along with the technological advancements. It was shown that such a system could be constructed with a BER of less than 10^{-15} by use of current research level technology, and it is expected that the bandwidth, latency, and scalability will directly scale with advances in the tunable VCSEL technology. An overview of a proposed extension of the optical crossbar for interconnecting multiple clusters was also presented. By combining WDM with space division multiplexing, it should be possible to extend the crossbar network to systems containing a much larger number of processors. We plan to analyze this further in later research.

References

1. J. L. Jewell, Y. H. Lee, A. Scherer, S. L. McCall, J. P. Harbison, and L. T. Florez, "Surface-emitting microlasers for photonic switching and interchip connections," Opt. Eng. **29,** 210–214 (1990).
2. R. A. Morgan, "Advances in vertical-cavity surface-emitting lasers," in *Vertical-Cavity Surface-Emitting Laser Arrays*, J. L. Jewell, ed., Proc. SPIE **2147,** 97–119 (1994).

3. J. Neff, "Optical interconnects based on two-dimensional VCSEL arrays," in *IEEE Proceedings of the First International Workshop on Massively Parallel Processing Using Optical Interconnections* (IEEE Computer Society, Los Alamitos, Calif., 1994), pp. 202–212.
4. F. Sugihwo, M. Larson, and J. S. Harris, Jr., "Low threshold continuously tunable vertical-cavity surface-emitting lasers with 19.1 nm wavelength range," Appl. Phys. Lett. **70,** 547–549 (1997).
5. M. Larson, F. Sugihwo, A. Massengale, and J. S. Harris, Jr., "Micromachined tunable vertical-cavity surface-emitting lasers," in *Proceedings of the International Electron Device Meetings* (IEEE Electron Devices Society, New York, 1996), pp. 405–408.
6. F. Sugihwo, M. Larson, C. C. Lin, W. Martin, and J. S. Harris, Jr., "25 nm wavelength range tunable vertical cavity lasers," in *Proceedings of Device Research Conference* (IEEE Electron Devices Society, New York, 1997), pp. 108–109.
7. F. Sugihwo, M. Larson, and J. S. Harris, Jr., "Monolithically micromachined wavelength tunable vertical cavity lasers," in *Proceedings of State of the Art on Compound Semiconductors XXVII* (Electrochemical Society Electronics Division, Paris, 1997), pp. 118–124.
8. M. Y. Li, W. Yuen, and C. J. Chang-Hasnain, "Top-emitting micromechanical VCSEL with a 31.6 nm tuning range," IEEE Photonics. Technol. Lett. **10,** 18–20 (1998).
9. C. A. Brackett, "Dense wavelength division multiplexing networks: principles and applications," IEEE J. Sel. Areas Commun. **8,** 948–964 (1990).
10. L. Kleinrock, M. Gerla, N. Bambos, J. Cong, E. Gafni, L. Bergman, J. Bannister, S. Monacos, T. Bujewski, and P.-C. Hu, "The supercomputer supernet testbed: WDM-based supercomputer interconnect," J. Lightwave Technol. **14,** 1388–1399 (1996).
11. P. Dowd, K. Bogineni, K. A. Aly, and J. A. Perreult, "Hierarchical scalable photonic architectures for high performance processor interconnection," IEEE Trans. Comput. **42**, 1105–1120 (1993).
12. K. M. Sivalingam and P. W. Dowd, "A multilevel WDM access protocol for an optically interconnected multiprocessor system," J. Lightwave Technol. **13,** 2152–2167 (1995).
13. E. R. Hedin and F. J. Goetz, "Experimental studies of electrooptic polymer modulators and waveguides," Appl. Opt. **34,** 1554–1561 (1995).
14. L. Robitaille, C. L. Callender, and J. P. Noad, "Design and fabrication of low-loss polymer waveguide components for onchip optical interconnection," IEEE Photonics Technol. Lett. **8,** 1647–1649 (1996).
15. Y. S. Liu, R. J. Wojnarowski, W. A. Hennessy, P. A. Piacente, J. Rowlette, Jr., M. Kadar-Kallen, J. Stack, Yu. Liu, A. Peczalaski, A. Nahata, and J. Yardley, "Plastic VCSEL array packaging and high density polymer waveguides for board and backplane optical interconnects," in *Proceedings of the 1998 IEEE Electronic Components and Technologies Conference* (Institute of Electrical and Electronics Engineers, New York, 1998), pp. 999–1005.
16. L. A. Hornak, ed., *Polymers For Lightwave and Integrated Optics* (Marcel Dekker, New York, 1992).
17. A. Louri and R. Gupta, "Hierarchical optical ring interconnection (HORN): scalable interconnection network for multiprocessors and multicomputers," Appl. Opt. **36,** 430–442 (1997).
18. A. Louri, B. Weech, and C. Neocleous, "A spanning multichannel linked hypercube: a gradually scalable optical interconnection network for massively parallel computing," IEEE Trans. Parallel Distribut. Syst. **9,** 497–512 (1998).
19. E. C. Vail, G. S. Li, W. Yuen, and C. J. Chang-Hasnain, "High performance and novel effects of micromachined tunable vertical-cavity lasers," IEEE J. Sel. Top. Quantum Electron. **3,** 691–697 (1997).
20. P. Tayebati, P. Wang, D. Vakhshoori, C.-C. Lu, M. Azimi, and R. N. Sacks, "Half-symmetric cavity tunable microelectromechanical VCSEL with single spatial mode," IEEE Photonics Technol. Lett. **10,** 1679–1681 (1998).
21. J. Chang-Hasnain, J. P. Harbison, C.-E. Zah, M. W. Maeda, L. Florez, N. Stoffel, and T.-P. Lee, "Multiple wavelength tunable surface-emitting laser arrays," IEEE J. Quantum Electron. **27,** 1368–1376 (1991).
22. H. Saito, I. Ogura, Y. Sugimoto, and K. Kasahara, "Monolithic integration of multiple wavelength vertical-cavity surface-emitting lasers by mask molecular beam epitaxy," Appl. Phys. Lett. **66,** 2466–2468 (1995).
23. Y. Li, T. Wang, and R. A. Linke, "VCSEL-based angle-multiplexed optoelectronic crossbar interconnects," Appl. Opt. **35,** 1282–1295 (1996).
24. R. Pu, E. M. Hayes, C. W. Wilmsen, K. D. Ohoquette, H. Q. Hou, and K. M. Geib, "Comparison of techniques for bonding VCSELs directly to ICs," J. Opt. A: Pure Appl. Opt **1,** 324–329 (1999).
25. H.-J. J. Yeh and J. S. Smith, "Integration of GaAs vertical-cavity surface emitting laser on Si by substrate removal," Appl. Phys. Lett. **64,** 1466–1468 (1994).
26. L. Eldada, C. C. Xu, K. M. T. Stengel, L. W. Shacklette, and J. T. Yardley, "Laser-fabricated low-loss single-mode raised-rib waveguiding devices in polymers," J. Lightwave Technol. **14,** 1704–1713 (1996).
27. H. J. R. Dutton, *Understanding Optical Communications* (Prentice Hall, Upper Saddle River, N.J., 1998).
28. M. C. Hutley, *Diffraction Gratings* (Academic, London, 1982).
29. J. W. Goodman, *Introduction to Fourier Optics* (McGraw-Hill, New York, 1996).
30. F. Timofeev, P. Bayvel, J. E. Midwinter, and M. Sokolskii, "High-performance, free-space ruled concave grating demultiplexer," Electron. Lett. **31,** 2200–2201 (1995).
31. T. V. Moui, "Receiver design for high-speed optical fiber systems," IEEE J. Lightwave Technol. **LT-2,** 243–267 (1984).
32. L. W. Shacklette, K. M. T. Stengel, C. Xu, and J. T. Yardley, "Polymeric waveguides for optical backplanes," in *Fly-by-Light: Technology Transfer*, D. B. Thompson, R. J. Baumbick, and L. B. Stotts, eds., Proc. SPIE **2467,** 107–117 (1995).