High-Level Model of a WDMA Passive Optical Bus for a Reconfigurable Multiprocessor System

V. E. Boros
Department of Computer Science and Electrical Engineering
The University of Queensland
St Lucia, Brisbane, Qld, 4072, Australia
borosv@csee.uq.edu.au

A. D. Rakić
Department of Computer Science and Electrical Engineering
The University of Queensland
St Lucia, Brisbane, Qld, 4072, Australia
rakic@csee.uq.edu.au

S. Parameswaran
Department of Computer Science and Electrical Engineering
The University of Queensland
St Lucia, Brisbane, Qld, 4072, Australia
sridevan@csee.uq.edu.au

ABSTRACT

We describe the first iteration of a comprehensive model with which we can investigate the practical limits on optical bus bandwidth and number of bus processing modules for given signal power. The selection algorithm will ultimately allow programmable evaluation of system parameters bus bandwidth, optical power budget, electrical power budget, number of modules and space consumption for an optimal design that is suitable for on-the-fly system reconfiguration.

Categories and Subject Descriptors

B.4.3 [**Input/Output and Data Communications**]: Interconnections (Subsystems)—*fiber optics, topology*

General Terms

Algorithms, Design

1. INTRODUCTION

In this paper we set out to describe a new interconnection model for multiprocessor computers. The model is for an optical bus comprising passive components. These components are collimating lenses, waveguide couplers and splitters, broadband beamsamplers, and filters. The filters accomplish an addressing function through wavelength division multiple access (WDMA) (and optionally through space multiplexing). Each system module has assigned to it a wavelength, which it uses to communicate with all other modules. On-the-fly system reconfiguration (e.g. load balancing) is also viable due to the greatly reduced intermodule communication latency compared with electrical connections. All communications share the same physical space on the bus. Intersignal interference is eliminated on the bus by wavelength separation (and optionally at the detectors by spatial separation). Our initial model allows for 128-bit electrical data and control line signals from a system module to be time division multiplexed (TDM) with an external modulator onto a single laser beam at an effective rate of 1.25 GHz (160 GHz serial bit rate[1]). Alternatively, an array of laser diodes can be used for parallel transmission, internally modulated at a few gigahertz. (Address line signals are converted to a bit sequence at the carrier wavelength ahead of the data and control signals.) This falls well within the fundamental frequency limit of light and the practical dispersive, attenuative limits of optical materials. The initial iteration of our model will also permit a bus that can, depending on the number of wavelength windows available, reach a few hundred modules before optical signal regeneration is required.

1.1 Motivation

Justification for research into optical interconnections can be found in the fractional bus speed/processor speed ratios of current commercial machines, the performance overheads of electrical interconnections in large scale multiprocessor systems, and the inherently large bandwidth available from light (hundreds of terahertz).

Bus access arbitration mechanisms required in an electrical bus are eliminated in the optical bus. Signals from all modules can be on the optical bus concurrently in time, and concurrently in space. This is possible because light does not interact with light directly.[2]

[1]Achieving this modulation rate on a single laser diode is a practical issue not addressed in this paper. This is not theoretically a problem; several lower rate modulators can be cascaded, staggered in space and thus in time to produce the required modulation rate in aggregate [14]. However the cost may be prohibitive.

[2]Light does interact readily with active media, but the optical crosstalk arising from this can be eliminated through basic design (wavelength filtering, spatial filtering). Light can also interact to some extent with passive media, giving rise to mode coupling, but this also can be minimized or eliminated through design.

1.2 Related Research

Related research over the past several years has yielded a point-to-point single hop network using passive components and space division multiplexing (SDM) for address selection [1], a multihop bus using passive components, active components and WDMA for address selection [11, 10], and a series of point-to-point interconnections that rely on an array of lasers (SDM) and WDMA to permit simultaneous and reconfigurable communication [15, 9]. Reconfiguration was achieved with tunable lasers and detectors the opacity or transparency of which depended on the incident wavelength. We seek to go beyond point-to-point and multihop networks, and to extend the concept of a single hop optical bus that permits reconfiguration of system level communication through WDMA.

There has been interest in the development of the partitioned optical passive star (POPS) interconnection network [5, 4]. This nonhierarchical network topology uses time division multiplexing (TDM) for multiple transmitter links to share star couplers. WDMA was rejected due to the required number of wavelengths needed to uniquely address each possible transmitter/receiver link. At the time 1 nm wavelength separations were available, severely limiting the number of wavelengths that could be transmitted over a fiber link. An alternative was wavelength reuse, but the need to reuse a wavelength to address another module required tuning times for lasers and detectors significantly less than the several nanoseconds and several microseconds that were available.

In the POPS topology, the star couplers restrict usage to one transmitting module per time slot in the case of TDM, but will permit different wavelengths at a time if using WDMA. In developing the optical eqivalent of a computer's system bus, we incorporate WDMA to permit multiple wavelengths concurrently in both time and space. One wavelength is assigned per module. The wavelength is used to identify a transmitter. To address the receiver it is necessary to incorporate a code at the start of the transmission. Broadcasting would require its own code. Bus access arbitration is eliminated. We believe that inclusion of coding constitutes minimal overhead, while permitting the number of wavelengths required for N modules to be N rather than N^2.

In regard to tunable lasers for extensive WDMA, much progress has been made over the last few years, for instance the development of widely tunable micromechanical lasers in which the dimensions of the resonant cavity are altered [2] to give a continuous tuning range of 32 nm and linewidth of about 0.5 nm. Directly tunable filters have also been developed [8], giving tuning ranges up to 62 nm and linewidths as small as 0.6 nm and aborption efficiency of 20%, with linewidths less than 1 nm over a 40 nm range. Smaller linewidths have also been reported [6] — 0.27 nm and a tuning range of 32 nm with 5.5 period TiO_2/SiO_2 Bragg mirrors; however analysis of sensitivity to tilt angle showed that 0.02° tilt resulted in an absorption drop of 18.2 dB.

Monolithic arrays of integrated lasers and Resonant Cavity Enhanced PhotoDetectors (RCEPDs) have been reported [3] in which wavelength tuning is accomplished over a spatial range. RCEPDs are grown beside the corresponding lasers to produce a very close wavelength match. By selectively enhancing and reducing the metal organic chemical vapour deposition (MOCVD) rate locally on the one substrate, an array of laser-RCEPD pairs, each with a different resonant wavelength, can be produced. The laser linewidths are small fractions of a nanometer. The RCEPD linewidths are a few tenths of a nanometer for 20% absorption efficiency, but increase to 2 nm in InGaAs for absorption around 80%.

In the following sections we seek to demonstrate the feasibility of a single hop optical bus using passive components and WDMA for addressing and system reconfiguration.

2. THE OPTICAL BUS BASED ON WDMA AND PASSIVE COMPONENTS

We show in Fig. 1 a framework for 1 master processing module, 3 slave processing modules and 1 shared (global) memory module. To support intermodule communication, each module includes a laser diode (LD) (alternatively, a laser array for reduced serial bandwidth) and driver, 1 photodetector (PD) array and receiver array, and LD modulators and PD demodulators. Each module is associated with a filter, which performs the wavelength addressing function. Each slave processing module is associated with a pair of broadband beamsamplers and a waveguide coupler. Depending on the amount of reflection at the air-beamsampler interface (beamsampling), attenuation and number of available wavelength windows, hundreds of slave processing modules can be inserted between the third slave processing module and the shared memory module.

Each module has assigned to it a transmit wavelength. Any module needing to communicate with a particular module transmits its signal, its wavelength identifying the sender at the receiver. Thus modules 0, 1, 2, 3, 4 are assigned wavelengths λ_0, λ_1, λ_2, λ_3, λ_4. Each filter is "tuned" such that only the wavelength assigned to a specific photodetector can reach that detector. To identify the required destination, the sending module adds a unique bit sequence at the commencement of the communication. Module 4, in sending data to module 3, produces a beam from its laser at λ_4 with a unique bit sequence identifying the intended destination. The λ_4 signal from module 4 will only reach detector 4 in each module. Signal detection at PD 4 thus identifies the sender to be module 4, and the initial bit sequence identifies the intended recipient as being module 3. Thus module 3 accepts the transmission from module 4, whereas the other modules ignore it.

The bus comprises two one-directional optical paths. Tapping beams off a path is achieved with the beamsamplers. Coupling beams into a path is achieved with waveguide couplers. For example, to couple the signal from module 3 so that it will reach module 4, the beam from module 3 is launched directly into a waveguide. Half of this beam (split by the coupler) is then coupled directly into the left going optical path. To couple the signal from module 3 so that it will reach module 1, after splitting by the waveguide coupler it is coupled directly into the right going optical path.

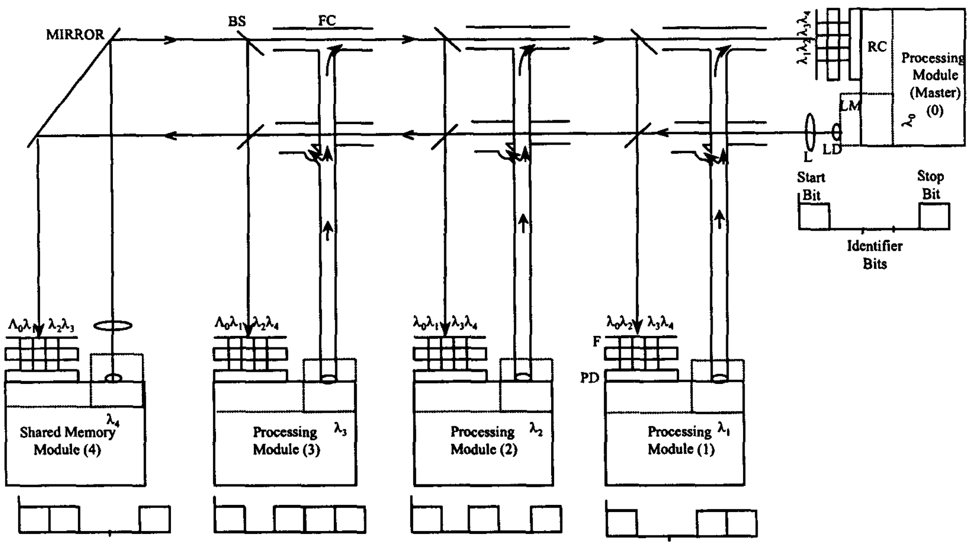

Figure 1: Optical bus concept based on WDMA. F = filter, PD = photedetectors, L = imaging lens, LM = laser modulators, LD = laser diodes, FC = fiber coupler, BS = beamsampler, RC = receiver circuit

Electrical address lines in each module determine the initial bit sequence for transmission, thus identifying the intended recipient. Electrical data and control lines are fed to the LD modulators and drawn from the PD demodulators.

The electrically decoupled signals permit light speed communication between modules. This can be used for efficient reconfiguration of the system, for example by shifting code or data from one module to another for load balancing.

We have attempted to preserve modularity in our design to permit extension to highly parallel systems. Thus we have restricted the function of our waveguide couplers to ramping a signal onto a free space optical path. Our design presently permits the extension of our optical bus to many modules in a line, rather than the 4 shown in Fig. 1, depending on power and wavelength considerations. The free space links will allow the design itself to be extended to permit many such lines in a row, many such lines in an orthogonal row in the same plane, and then rows orthogonal to the plane.

2.1 System Parameters Requiring Analysis

The system parameters for the design comprise intermodule data transfer rate, wavelength availability, optical power consumption, electrical power consumption, space consumption and system extendibility.

Intermodule Data Transfer Rate

The fundamental limit in the optical data rate for one laser beam is the frequency of light, in the order of hundreds of terahertz. The practical limit arises from the broadening of light pulses in optical media and loss of synchonisation in clocked buses due to dispersion mechanisms. The device parameters requiring analysis for pulse broadening comprise intermodal dispersion, intramodal dispersion and optical component thicknesses. A further complication arises from optical crosstalk at the detectors and in the filters, resulting in signal noise. This needs to be addressed during system design, since it will affect the bit error rate (BER).

Wavelength Availability

The tuning range and linewidth of a laser technology may place limits on the number of wavelength windows available for signals. The 2 nm filter linewidths limit the number of windows to about 20 in the InGaAs filters of Ref. [3]. Reducing the linewidths to less than 0.3 nm, as in Ref. [6], produces on the order of 10 times more windows. This unfortunately introduces another complication in that the laser and filter spectra must be very closely matched. If they were to mismatch by only a few tenths of a nanometer, the communication link would be lost. Thus the slight wavelength drift observed in Vertical-Cavity Surface-Emitting Lasers (VCSELs) would become significant, requiring temperature or current control of the laser.

A way to increase wavelength availability is to incorporate wavelength reuse. This would greatly reduce the effects of wavelength drift in lasers by permitting filters with large absorption windows, but would require additional methods of uniquely identifying each module. SDM could be used to accomplish this. In this regard the integrated VCSEL-RCEPD arrays of Ref. [3] would be useful.

Optical Power Consumption

The parameters for optical power consumption comprise absorption and scattering in optical media (beam attenuation), the amount of reflection at the air-beamsampler interfaces, and the absorption efficiency of the filters.

Electrical Power Consumption
Electrical power parameters comprise power requirements for the electrically isolated multi-chip modules, which include the LD drivers and modulators and the PD receivers and demodulators. We also need to account for losses from electro-optical signal conversion.

Space Consumption
The optical bus requires space for the drivers, receivers, modulators, demodulators, LDs, PDs and optical components. These replace the electrical bus and bus arbiter.

System extendibility
The preceding parameters, once investigated, can be used to develop a parameterized model for system extendibility, relating available optical and electrical power, minimum required data transfer rate, available wavelength windows and available space to maximum number of modules.

3. HIGH-LEVEL MODEL

In this first model we do not consider the electrical power and space analysis, but focus on the limits of optical data transfer speed due to dispersion and number of interconnectable modules due to available optical power and wavelength windows. Our present 160 GHz configuration is well within the practical dispersive limits of optical media, since a bandwidth-length product for single mode step index optical fibres of 100 GHz km is attainable. Since this modulation frequency requires a cascade of expensive external laser modulators, we could incorporate SDM and use arrays of (gigahertz) internally modulated lasers. The models we are beginning to develop will determine finally the maximum aggregate bus bandwidth possible for our system, and the maximum number of modules that can be reached before optical signal regeneration is required, subject to wavelength availability.

3.1 Optical Pulse Broadening and Desynchronization

3.1.1 Intermodal Dispersion

Surface emitting lasers usually operate in several modes (with different frequencies and different light intensity distributions) simultaneously. The optical field will see different refractive indices depending on its modes of propagation, resulting in intermodal dispersion and thus pulse broadening in multimode pulses (or desynchronization in clocked buses that consider only the free space light delay).

To find the rms pulse broadening, we calculate the time delay difference of the slowest and fastest modes in a section of length L [12]:

$$\tau_{\text{mod}} = L(\bar{n}_{\max} - \bar{n}_{\min})/c = L\delta\bar{n}/c \quad [\text{s}]. \tag{1}$$

We consider the power in all the modes as being normalised:

$$\sum_{i=1}^{q} P(t_i) = 1, \tag{2}$$

where $P(t_i)$ indicates the proportion of power in existence at time t_i. By centering about $t = 0$ and using Eq. (1), it should be clear that the fastest mode experiences delay $t_1 = -L\delta\bar{n}/(2c)$ whereas the slowest mode experiences delay $t_q = L\delta\bar{n}/(2c)$.

The rms pulse broadening is given by the standard deviation σ_{mod} of the t_i [13]:

$$\sigma_{\text{mod}} = \sqrt{\sum_{i=1}^{q}(t_i P(t_i))^2 - \left(\sum_{i=1}^{q} t_i P(t_i)\right)^2} \quad [\text{s}]. \tag{3}$$

3.1.2 Intramodal Dispersion

Intramodal dispersion mechanisms comprise chromatic dispersion and polarization dispersion. For this first iteration of the model we ignore the effects of the latter.

3.1.2.1 Chromatic Dispersion

Chromatic dispersion arises from the wavelength dependence of the refractive index. This is theoretically an issue, since light sources are not purely monochromatic.

The rms pulse broadening σ_c from a source having spectral width σ_λ is given by [7]

$$\sigma_c \simeq \sigma_\lambda L D_C \simeq \sigma_\lambda L \frac{-2\pi}{c\lambda^2} \frac{d\bar{n}_g}{dk_0} \quad [\text{s}], \tag{4}$$

where L is the thickness or length of the optical medium, c the free space speed of light, λ the source central wavelength, and $\bar{n}_g$ the group mode index.

The derivatives of $\bar{n}$ can be ascertained by numerical techniques from the eigenvalue equation.

3.2 A Model for Intermodule Data Transfer Rate

In developing an initial model for the best data transfer rate, we disregard optical crosstalk. We are in process of developing for our system a crosstalk model that can minimize signal noise (hence BER) due to a combination of wavelength and spatial filtering.

By assuming that the temporal evolution of the input pulse can be considered random with a gaussian (normal) distribution, the 3 dB optical bandwidth limited by intermodal dispersion is [12]

$$B_{\text{mod}}^{(\text{opt})} = 0.187 / \sum_{j=1}^{r} \sigma_{\text{mod}_j} \quad [\text{Hz}], \tag{5}$$

where r is the number of components in the optical path. Similarly the 3 dB optical bandwidth limited by chromatic dispersion is

$$B_c^{(\text{opt})} = 0.187 / \sum_{j=1}^{r} \sigma_{c_j} \quad [\text{Hz}]. \tag{6}$$

Overall bandwidth will be a function of $B_{\text{mod}}^{(\text{opt})}$ and $B_c^{(\text{opt})}$.

To illustrate how these equations are used to calculate the optical bandwidth, we will assume that the bus is asynchronous and that the optical pulses are single mode. Thus we use Eq. (4) to calculate the pulse broadening from chromatic dispersion, and Eq. (6) to calculate the resultant bandwidth. To find the chromatic dispersion we can evaluate $d\bar{n}_g/dk_0$ numerically from the eigenvalue equation. We perform these calculations first for the beamsamplers, second for the beam couplers, and then add the results, using Eq. (5) to calculate the resultant modal bandwidth.

3.3 A Model for Optical Power Consumption

In developing an initial optical power budget, we neglect the effects of attenuation in optical media and assume lossless materials. In practice the attenuation from absorption of light in BK7 glass (used in our broadband beamsamplers) is minimal. We neglect any reflection from antireflection coated surfaces. We assume that all light falling on a photodetector is incident only on the active area.

From Fig. 1, the longest path with regard to reflection loss (from beamsamplers) and splitting loss (from waveguide 3 dB power splitters) is from module 2 to module $(l-1)$ or conversely. The minimum power available at the PD is

$$P_{\mathrm{D}} = 0.5 T_{\mathrm{F}} R_{\mathrm{BS}} T_{\mathrm{BS}}^{l-3} P_{\mathrm{L}} \quad l = 4, 5, \ldots \quad [\mathrm{W}], \qquad (7)$$

where l is the number of modules, T is the intensity transmission coefficient, R the intensity reflection coefficient, F means *filter* and BS means *beamsampler*.

The laser's spectral emission shape needs to be integrated over the wavelength range that falls inside the filter window. The proportion of this to the total integral defines the percentage of light incident on the filter that is allowed to transmit. Thus T_{F} is given by

$$T_{\mathrm{F}} = \left(\int T(\lambda)P(\lambda)\mathrm{d}\lambda\right) / \left(\int P(\lambda)\mathrm{d}\lambda\right), \qquad (8)$$

where P is the spectral light power density at wavelength λ and T is the intensity transmission coefficient of the filter window, also at λ.

Eq. (7) is the key to determining required optical power levels given a specified number of modules, or conversely, calculating maximum allowed modules given optical power restrictions.

To illustrate the application of Eq. (7), let us restrict the available optical power from the laser diode to 1 mW, and use photodetectors having a noise equivalent power (NEP) rating of 900 $\mathrm{pW}/\sqrt{\mathrm{Hz}}$. 0.1% of the light incident on a beamsampler is reflected toward the filter. We use Eq. (8) along with the filter transmission characteristics and the LD emission characteristics to calculate the amount of light allowed through the filter. For the purpose of this illustration we will assume that 1% of light incident on a filter is transmitted. Solving Eq. (7) for l yields 331 modules that can be attached to the bus. This requires 331 wavelength windows, a difficult achievement in any current technology, unless we were to incorporate wavelength reuse and SDM.

3.4 Wavelength Availability and SDM

The number of wavelength windows, W, available in a laser-filter-detector (LFD) combination for a chosen LFD technology and wavelength control technology must satisfy the targeted BER. If W is less than the required number of modules, we can exploit spatial parallelism and reuse wavelength windows in spatially adjacent paths. Because spatial information is lost inside fibers, we need also to exploit spatial parallelism on the bus paths. For a given number of lasers, N, and the number of wavelength windows, W, meeting the BER requirement, the number of uniquely identifiable modules, l, is given by

$$l = W \times N \qquad (9)$$

4. A SELECTION ALGORITHM

We present an algorithm for determining up to three system parameters, based on the models we have discussed. These parameters are bus bandwidth, optical output power and number of modules. Specifications involving these parameters will stipulate some combination of minimum required aggregate bandwidth `minABW`, maximum available optical output power `maxOP`, maximum available number of wavelength windows `maxW`, minimum required number of modules `minM` and maximum allowed bite error rate `maxBER`. Solutions will be a combination of maximum allowed number of modules `maxM`, minimum required optical output power `minOP`, and maximum achievable aggregate bandwidth `maxABW`. Other parameters are total pulse broadening `totalPB`, module pulse broadening `modulePB`, dispersion limited bandwidth `dispBW`, modulation limited bandwidth `modBW`, minimum and maximum number of laser-filter-detector combinations `minLFD` and `maxLFD`, and error functions $\mathcal{E}$ for different combinations of `minBW`, `maxOP`, `minM`.

```
find maxW for maxBER and current WDM and wavelength
control technologies
find modBW for maxBER and chosen modulation
technology
  if maxOP then
    find maxM for maxOP:
      maximize l (Eq. (7), Eq. (9): N = maxOP/P_L)
      maxM = maximized l
      maxLFD = N
    find maxABW for maxM and maxLFD:
      calculate modulePB (Eqs. (3), (4))
      totalPB = modulePB × maxM
      calculate dispBW for totalPB and maxBER
(Eqs. (5), (6))
      maxABW = minimum[dispBW, modBW] × maxLFD
  if minM then
    find minOP for minM:
      minimize N × P_L (Eq. (7), Eq. (9): l = minM)
      minOP = minimized N × P_L
      minLFD = minimized N
    find maxABW for minM and minLFD:
      calculate modulePB (Eqs. (3), (4))
      totalPB = modulePB × minM
      calculate dispBW for totalPB (Eqs. (5), (6))
      maxABW = minimum[dispBW, modBW] × minLFD
  if minABW then
    minLFD = minABW/modBW
    find maxM for minABW:
      calculate modulePB (Eqs. (3), (4))
      calculate totalPB for modBW (Eqs. (5), (6))
      maxM = ⌊totalPB/modulePB⌋
```

```
    find minOP for maxM and minLFD
if (maxOP and minABW) then
    find minimum[maxM for maxOP, maxM for minABW]
if (maxOP and minM) then
    find maxM for maxOP
    if maxM ≥ minM then
        null
    else
        minimize ℰ(maxOP, minM)
if (minM and minABW) then
    find maxM for minABW
    if maxM ≥ minM then
        null
    else
        minimize ℰ(minM, minABW)
if (maxOP and minM and minABW) then
    find minimum[maxM for maxOP, maxM for minABW]
    if maxM ≥ minM then
        null
    else
        minimize ℰ(maxOP, minM, minABW)
```

With regard to finding solutions for Eq. (7) in the algorithm, we note that in this model T_F and P_D are constant.

The error functions are required for specifications in which 2 or more of the specified parameters cannot all be satisfied in the one design. We make $\mathcal{E}$ to be a function of weighted parameter errors, the weightings an indication of the importance attached to the error.

Using this algorithm, we are able to specify an optical bus based on our design that will meet user requirements of minimum required optical bandwidth, maximum available optical output power and minimum required number of modules.

5. CONCLUSION

We have presented a new scheme for optical communication among modules of a multiprocessor computer system. This scheme uses WDMA and tunable lasers and filters to perform an addressing function and to allow efficient system reconfiguration. The interconnection network comprises passive optical components — beamsamplers, waveguide couplers and filters, permitting single hop communication between modules. This permits at least 128 electrical signals at 1.25 GHz to be modulated onto a single optical bus to a few hundred modules. We have also presented the first iteration of a comprehensive model and selection algorithm that will allow optimal design for any combination of bus bandwidth, optical power budget, electrical power budget, number of available wavelengths, number of system modules and space consumption.

6. REFERENCES

[1] S. Araki, M. Kajita, K. Kasahara, K. Kubota, K. Kurihara, I. Redmond, E. Schenfeld, and T. Suzaki. Experimental free-space optical network for massively parallel computers. *Appl. Opt.*, 35(8):1269–1281, 10 Mar. 1996.

[2] C. J. Chang-Hasnain. Widely tunable micromechanical VCSEL. In *Proc. IPC'98*, pages 116–117, 15–18 Dec. 1998.

[3] J. Cheng and Y. Zhou. A new wave in network data links. *IEEE Circuits and Devices*, pages 17–27, Nov. 1999.

[4] D. M. Chiarulli, S. P. Levitan, R. P. Melhem, J. P. Teza, and G. Gravenstreter. Partitioned optical passive star (POPS) multiprocessor interconnection networks with distributed control. *J. Lightwave Technol.*, 14(7):1601–1612, Jul. 1996.

[5] G. Gravenstreter, R. G. Melhem, D. M. Chiarulli, S. P. Levitan, and J. P. Teza. The partitioned optical passive stars (POPS) topology. In *Proc. 9th Int. Parallel Process. Symp.*, pages 4–10. IEEE Comput. Soc. Press, 1995.

[6] H. L. Hartnagel, J. Pfeiffer, K. Mutamba, J. Peerlings, R. Riemenschneider, and P. Meissner. III-V membrane structures for tunable Fabry-Pérot filters and sensor applications. In *Proc. COMMAD'98*, pages 49–56. IEEE, 14–16 Dec. 1998.

[7] G. Keiser. *Optical Fiber Communications.* McGraw-Hill, 2nd ed., 1991.

[8] R. Le Dantec, T. Benyattou, G. Guillot, A. Spisser, C. Seassal, J. L. Leclercq, P. Viktorovitch, D. Rondi, and R. Blondeau. Tunable microcavity based on InP-air bragg mirrors. *IEEE J. Select. Topics Quantum Electron.*, 5(1):111–114, Jan./Feb. 1999.

[9] J. E. Leight, S. Homan, A. E. Willner, G. Giaretta, M. Y. Li, and C. J. Chang-Hasnain. Experimental demonstration of reconfigurable and simultaneous wavelength-division-multiplexed multiple-plane optical interconnections. *IEEE Photon. Technol. Lett.*, 8(2):302–304, Feb. 1996.

[10] Y. Li, A. W. Lohmann, Z. G. Pan, S. B. Rao, I. Redmond, and T. Wang. Optical multiple-access mesh-connected bus interconnects. *Proc. IEEE*, 82(11):1690–1700, Nov. 1994.

[11] Y. Li, A. W. Lohmann, and S. B. Rao. Free-space optical mesh-connected bus networks using wavelength-division multiple access. *Appl. Opt.*, 32(32):6425–6437, 10 Nov. 1993.

[12] J. M. Senior. *Optical Fiber Communications: Principles and Practice.* Prentice Hall, 2nd ed., 1992.

[13] R. E. Walpole and R. H. Meyers. *Probability and Statistics for Engineers and Scientists.* Macmillan Publ. Co., 5th ed., 1993.

[14] A. E. Willner. Mining the optical bandwidth for a terabit per second. *IEEE Spectr.*, 34(4):32–41, Apr. 1997.

[15] A. E. Willner, C. J. Chang-Hasnain, and J. E. Leight. 2-D WDM optical interconnections using multiple-wavelength VCSEL's for simultaneous and reconfigurable communication among many planes. *IEEE Photon. Technol. Lett.*, 5(7):838–841, Jul. 1993.