Channel Allocation, Power Budget And Bit Error Rate In Hierarchical Optical Ring Interconnection Network (HORN)

Thomas S. Jones and Ahmed Louri
Department of Electrical and Computer Engineering
The University of Arizona
Tucson, AZ 85721
louri@ece.arizona.edu

Abstract

Recent advances in computer processing speeds have resulted in a parallel processing environment in which the interconnection networks (INs) themselves are the limiting factor in terms of performance. Larger and faster INs can be implemented optically subject to the current limitation in the number of wavelengths imposed by optical switch and filter technology if the INs are arranged hierarchically, i.e. if the processing nodes are arranged in clusters and these clusters are connected together special routing nodes at one or more higher levels. HORN [15] uses optically connected rings as the basic building blocks and connects these rings using an optical tree. Three key issues facing HORN are addressed in this paper: dynamic channel allocation (DCA), optical power budget (OPB) and bit error rate (BER). Four approaches to DCA are evaluated and a design trade-off is performed between them. The first two are taken from the literature [5] while the last two are proposed here and expoit the multiple paths available in a hierarchical network. The evaluation of OPB and BER shows that HORN is feasible and practical when optical amplification is used at the initial signal insertion point for transmissions at higher levels in the hierarchy.

1. Introduction

It's generally known and accepted that the speeds at which computers process information has increased almost two orders of magnitude in the last ten to fifteen years. Technological advances in other areas in parallel processing systems haven't been able to keep up with that increase in processing speed, reaching the point that interconnection networks (INs) themselves are now the limiting factor in the performance of parallel processing systems [15, 7]. The application of optical interconnects to INs holds the potential of meeting the higher speeds and throughput demanded by higher CPU speeds. In theory, thousands of optical signals can be multiplexed onto the same optical fiber with much lower crosstalk between them than in traditional, electronic interconnection networks, but in practice, the limitations imposed by transmitters, detectors, splitters, etc. permit only ten to twenty channels per fiber [1, 6].

One approach to overcome this limitation is to organize processing nodes hierarchically [3, 5, 18]. Nodes are interconnected in small clusters and these clusters are connected at a new level either through interface nodes or through an entirely new network superimposed on the small clusters. This process is repeated as many times as necessary to achieve a network of the desired number of processing nodes while not violating the 10-20 wavelength limit per cluster. This approach is also supported by conclusions by Goodman [7] and Dandamudi [2] that nodes in an IN engage in data transfers more frequently with nearby nodes than with more distant ones, at least in a well designed IN. These hierarchical networks can be classified by the topology of the basic cluster and the topology by which clusters are interconnected. Proposed topologies include a star of stars [13], a bus of buses [1], a hypercube of hypercubes [8] and a tree of rings [15, 16].

Hierarchical Optical Ring Interconnection Network (HORN) [15] is an all-optical IN in which the basic clusters are rings that are interconnected in a tree structure. Major objectives for HORN include non-blocking, collision-free communication; scalability; unity diameter; low latency; and efficient accommodation of variations in communication traffic flow. An example of HORN is shown in Figure 1 in which 234 processors are connected in a three layer hierarchy. $H(n, g)$ is used in this figure to identify individual groups of a hierarchical interconnection, where n refers to the level and g refers to a group at level n; $H(n)$ is used to identify all groups of a hierarchical interconnection at level n.

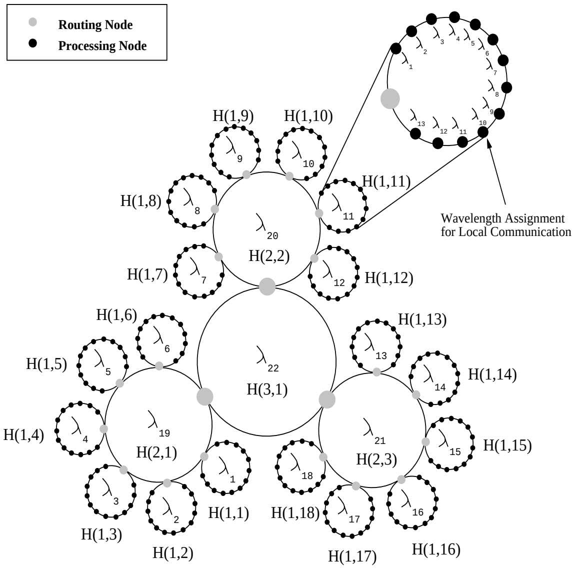

Figure 1. Sample HORN network with 234 processing nodes.

HORN uses wavelength division multiplexing (WDM) and spatial separation to isolate communication on the rings at different levels. On each of the lowest level rings, also called local rings, each processing node (PN) is assigned a unique wavelength for reception and a source selects a destination by transmitting on the wavelength assigned to the destination. The same set of wavelengths can be reused in each local ring since they are isolated from one another by the routing nodes. At the remote levels a wavelength is uniquely assigned to each *ring* of PNs, or ring of rings of PNs, depending on the level at which communication occurs. Wavelengths can also be reused at these levels as long as the isolation from the other rings in the network is maintained. It should be apparent that broadcasting and multi-casting is inherent in this approach; messages can be sent to all PNs on a ring by transmitting on the appropriate channel and level. In Figure 1 the same set of wavelengths that are shown for $H(1,11)$ is repeated in each of the local rings (any ring designated $H(1,g)$) thus permitting 256 virtual channels to operate simultaneously while using only 22 physical wavelengths.

Several key issues need to be addressed to establish the benefits and practicality of HORN. First, while WDM and spatial separation greatly reduce the number of contentions for each of the channels in question, there are still times in which two sources may attempt to access the same channel at the same time. A media access (MAC) protocol is needed to prevent this from happening and to ensure fairness of access to all sources. Second, data traffic is often bursty in parallel processing systems and this could result in some channels going unused in one portion of a IN while performance in another portion of the network is suffering due to overuse of the available channels. One approach to this problem that has been proposed is dynamic bandwidth reallocation [18, 5] wherein a system tracks some parameters in the network and dynamically reassigns channels to the areas with the greater needs. This reallocation approach, which we call dynamic channel allocation (DCA), should be analyzed with respect to need in HORN and a design/algorithm should be specified. Finally, the optical power budget (OPB) for HORN needs to be evaluated including signal to noise ratio (SNR) and bit error rate (BER).

An analysis MAC protocols was performed previously and a token based protocol was selected for use in HORN [12]. In this protocol, called THORN, different bits of a token packet are used to control access to each of the data channels. THORN ensures fairness of access to all nodes while also maximizing channel usage rate in that a node does not have to wait for an assigned slot to transmit its data. It also provides a method to transmit small, high priority messages virtually immediately. As will be shown later, THORN also provides another method to track channel usage for DCA.

This paper will analyze DCA and OPB. Two approaches to DCA are taken from the literature and two other approaches made possible by THORN are proposed. All four approaches to DCA will be analyzed and compared. Even though data transmission in HORN is “single hop”, signal losses are still proportional to the number of nodes that a message must pass through. An attempt is made in this paper to assess whether losses in HORN are low enough to ensure sufficient signal strength for reliable detection at every node in the network. The need for optical signal amplification is also analyzed. This analysis is based on the anticipated losses due to signal insertion, transmission, signal extraction using splitters at each of the PNs, and routing using acousto-optic tunable filters (AOTFs) to route messages between levels.

2. Analysis

2.1. Dynamic Channel Allocation

There are four possible approaches to DCA in HORN; the first two are presented in the literature [5] while the last two are made possible by the THORN protocol [12]. In the first, a complete representation of the pertinent communication parameters is maintained at every processing and routing node in the network. These parameters include channel assignments and some indication of current demands on the assigned channels. The second approach is similar but only a partial representation of the network is maintained at each node, data is maintained only for the current node and below in the hierarchy. In the third and fourth approaches each node monitors the tokens for each ring or PN directly connected to it at the next lower level in the hierarchy. For example, node (t) in Figure 2 monitors the tokens for nodes (a) through (f). In the first three approaches the channels are dynamically reassigned while the system is operating based on the demand indicated by the method chosen, but in the fourth approach communication is transferred to the next higher level in the hierarchy if the required channel(s) at the desired level are unavailable. These four approaches will be analyzed below in terms of database size, hardware complexity and communication overhead.

2.1.1 Full Database Representation Approach

In the full database representation approach the size of the database at each node will be proportional to the number of nodes in the system, which is given by:

$$N_T = \sum_{i=1}^{r} n_i \tag{1}$$

where n_i is the number of nodes at level i and r is the number of levels in the system. Since this database is duplicated at all nodes in the system the amount of data stored across all nodes is proportional to N_T^2. This approach requires multiple, tunable transmitters and receivers at each node since the physical wavelength assigned to a node will change based on availability and multiple transmissions will arrive simultaneously when multiple channels are assigned to a node at a given level. The software to reallocate channels and synchronize data transmission at different levels will be the most complicated of the four approaches since this reallocation, coordination and synchronization will have to be asserted across the entire network in concert. Any change in channel allocation will have to be transmitted simultaneously to every node in the network and the databases at each channel will have to be updated regularly even if no changes are made to channel allocations. This means that the communication traffic just to coordinate DCA will be proportional to N_T. Channel allocation will have to be synchronized across the entire network in order to maintain currency of the databases and to ensure that messages are not misdirected. Dowd, et. al. have concluded that the delay imposed to synchronize the data cycles under this approach is unacceptable [4].

2.1.2 Partial Database Representation Approach

In the partial database representation approach the size of the database at each node will be proportional to the number of nodes at the current level and below only, which is given by:

$$N_P = \sum_{i=1}^{j} n_i' \tag{2}$$

where n_i' is the number of nodes at level i in the local structure only and j is the number of the current level. The total size of the database across the entire system is proportional to $\Pi_{i=1}^{r} n_i \sum_{j=1}^{i} n_i'$. This approach also requires multiple, tunable transmitters and receivers at each node since the physical wavelength assigned to a node will change based on availability and multiple transmissions will arrive simultaneously when multiple channels are assigned to a node at a given level. The software to reallocate channels and synchronize data transmission at different levels will be the less complicated since this reallocation, coordination and synchronization will only have to be asserted across a portion of the network and any change in channel allocation only has to be transmitted to the affected nodes. Communication traffic to coordinate DCA will be proportional to N_P. Since channel allocation is localized synchronization of data cycles will also only affect portions of the network at different times. Dowd's comment that synchronization imposes unacceptable delays would especially apply to this approach since this is his favored approach [3, 5]. He has proposed an alternative that does not require data cycle synchronization [4].

2.1.3 Token Based Approach With Channel Reallocation

In this paper we propose two DCA approaches that monitor the tokens used by the THORN protocol. These tokens are a direct indication of channel usage and demand. In the first token based approach channels are reallocated based on the usage and demand indicated by the tokens. In the second approach, discussed more below, transmissions are redirected to higher levels in the hierarchy when the tokens indicate that the desired channel will not be free soon. Neither of these approaches use databases either at each node or at some central location. The token based approach with channel reallocation will require multiple, tunable transmitters and receivers at each node for each level of the hierarchy as will any approach in which channels are reassigned dynamically. The software to reallocate channels and synchronize data transmission will be even less complicated than in the partial database approach since this reallocation, coordination and synchronization will only have to be asserted between adjacent levels. Synchronization is greatly simplified since it can be controlled through the tokens. The THORN protocol presented previously [12] provides for a small payload as part of the token packet which can be used to coordinate and reassign channels. Communication traffic to coordinate DCA will therefore have minimal affect on data transmissions since the token packets circulate on dedicated channels.

2.1.4 Token Based Approach With Message Redirection

This last approach imposes minimal requirements on the types and number of transmitters and receivers at each node. The software will be very simple; all it has to do is determine if the desired channel is busy and redirect transmission to another level if it is. No synchronization of data cycles will be necessary and there is no communication overhead for DCA. In this approach the small payload in the token packets under THORN can be used to transmit small, high priority messages or even normal priority messages on a limited basis. This approach doesn't use DCA in the same way the other approaches since the channels are not reassigned but does fill the same need and makes use of the same channels that DCA does without having to perform the channel reallocation dynamically. The main drawback to this approach is that it places more demand on the channels assigned to higher levels in the hierarchy. The solution to this is to assign an increasing number of channels to each level just as a fat tree data structure doubles the number of links for each step up the tree.

2.2. Optical Power Budget

Routing in HORN is accomplished on two levels: local and remote. Local communication takes place when source and destination PNs are physically on the same lowest (local) level ring. Each PN is assigned a different wavelength for reception, so routing is simply a matter of tuning to the right wavelength. The same set of wavelengths is reused in every local ring but these messages never leave their own local ring so there are no channel conflicts and no other routing needs to be performed on them. Messages transmitted in all other cases are optically routed by wavelength to the destination *ring* through at least two routing nodes. Figure 2 shows a sample of remote routing where the message must be optically routed at nodes (g), (u), (t), and (f). The desired PN on the destination ring is specified in a simple header. Each PN on the target ring processes the header and either buffers the message or discards it based on the destination address in the header.

Two major issues in designing a HORN network are the routing nodes technology and the splitter design. At the routing nodes the objective is to direct optical signal based on their wavelength and to be able to control that routing by varying some parameter in the router. Two optical technologies are mature and able to meet these requirements: AOTFs [10] and Fabry-Perot Filters (FPFs) [9]. AOTFs operate on the principle that optical signals can be directed in very predictable ways based on their interaction with acoustic waves. FPFs use resonant cavities that transmit only transmissions whose wavelengths are exact multiples of the length of the cavities. AOTFs were selected because they are tunable over a broader range of wavelengths and can select more frequencies simultaneously [10]. AOTFs also have very high inband transmission and very low crosstalk. Inband transmission better than 99% and crosstalk of less than -15 dB have been reported [10, 11, 19].

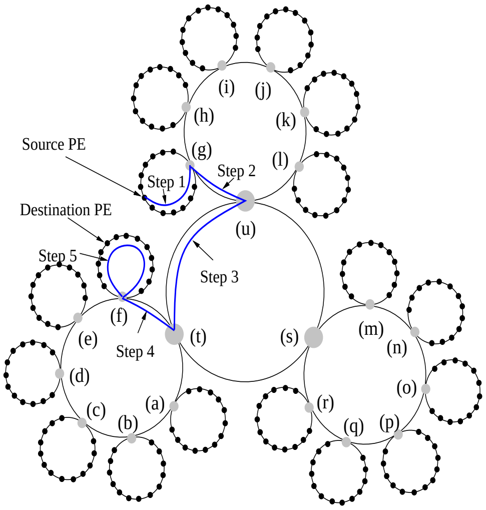

Figure 2. Routing of remote transmissions in HORN.

Message transmission is unidirectional on all rings so routing in both the local and remote cases can be modeled as a string of PN and routing nodes.

In the local transmission case a message will not have to pass through more than $N - 1$ splitters where N is the number of nodes on a local ring, as shown in Figure 3. Attenuation of the signal is the same at $N-2$ splitters but at the destination node the message is deflected to the PN rather than passed on the next node. There will be three distinct components to the loss: signal insertion loss, transmission loss at up to $N - 2$ splitters, and an extraction loss.

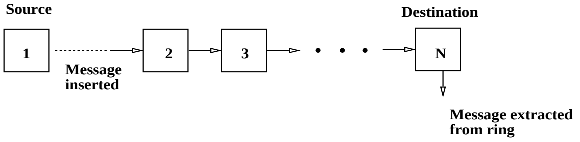

Figure 3. Equivalent model of a local message transmission path in HORN

Figure 4 shows the remote transmission case in which a message may have to pass through up to $2N-1$ splitters and up to N_r routing nodes (AOTFs) where N_r is the total number of routing nodes in the system. Signal loss in this case will be comprised of insertion loss, transmission loss through up to $2N-2$ splitters, loss at up to N_r AOTFs, and an extraction loss.

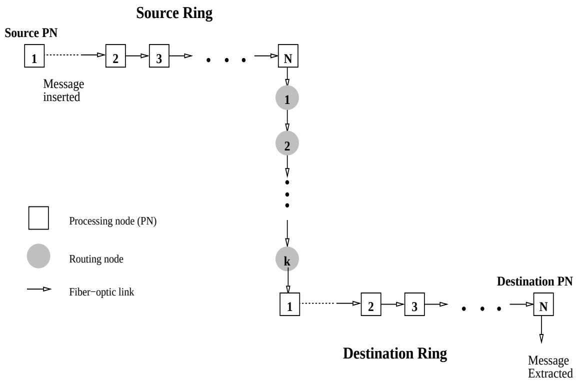

Figure 4. Equivalent model of a remote message transmission path in HORN

Fiber losses can be significant but losses less than 3 dB/km for $0.8\mu m \leq \lambda \leq 1.4\mu m$ are possible [14]. Since total fiber length will be on the order of 10 meters the fiber losses for both local and remote transmissions can be made to be negligibly small and will therefore be ignored in the following calculations

At the PNs the splitter selection is a balance between extracting enough signal for detection at each node and transmitting as much signal as possible to the next node. Splitters are available in a wide range of transmission efficiencies. A high transmission efficiency means that a low percentage of the power is extracted for use at the node in question so if the efficiency is too high then not enough power is being extracted to detect. On the other hand a low transmission efficiency means that too little power may be passed on to the rest of the nodes in the ring and the signal will die out after passing through only a few PNs.

Using the model given in Figure 3, an expression can be derived for the minimum power reaching the detector at a destination node for a local transmission when there are N PNs on each local ring:

$$P_{out} = P_{in} + L_i + R + (N-2)(T+L_c) \tag{3}$$

where P_{out} is the power reaching the detector, P_{in} is the strength of the source, L_i is the insertion loss, R is the percentage of the signal that is extracted or deflected at each splitter, T is the splitter transmission efficiency and L_c is the coupling efficiency between the fiber and the splitters. (All quantities are in dB.)

Similarly, using the model in Figure 4, an expression can be derived for the minimum power reaching a detector for remote transmissions in HORN when there are N PNs on each local ring and a total of N_r routing nodes in the system:

$$\begin{aligned} P_{out} &= P_{in} + L_i + R + (2N-2)(T+L_c) \\ &\quad + k(L_{AOTF} + L_c) \end{aligned} \tag{4}$$

where L_{AOTF} is the transmission loss through an AOTF. As stated previously, AOTF transmission efficiencies of more than 99% have been reported, corresponding to an L_{AOTF} of -0.043 dB at each AOTF [10, 11, 19].

Insertion losses of 1 dB are routinely achieved in optical fibers. Note that $10^{\frac{T}{10}} + 10^{\frac{R}{10}} = 1$ since the amount of power into the splitter has to be equal to the power transmitted and reflected from it, so the selection of a transmission efficiency defines both T and R. The coupling efficiency, L_c, is determined by four interfaces where the index of refraction changes, two at the two ends of the fibers and two at the input and output faces of the splitter. If air is in the spaces between the fibers and the splitter then it's possible to achieve a coupling efficiency of $L_c = 10\log 0.96^4 = -0.7$ dB since the reflectance for normal incidence at an air to glass interfaces is typically about 4%. If index matching oil is instead in those spaces then it should be possible to obtain $L_c = -0.5$ dB.

Using these results we can derive an expression for the maximum number of PNs as a function of splitter transmission efficiency for local transmissions:

$$N_{max} = \left(\frac{P_{out(min)} - P_{in} - L_i - R}{T + L_c} \right) + 2 \tag{5}$$

where $P_{out(min)}$ is the the minimum power for reliable detection.

Equation (5) is graphed in Figure 5 for $P_{in} = 1W$ and $P_{out(min)} = 10\mu W$. It can be seen from this figure that splitters with transmission efficiencies of 90% will allow a system with a maximum of 43 PNs. This far exceeds the system constraints in HORN of no more than about 20 PNs on a local ring but only exceeds the minimum requirement for remote transmissions by three PNs since data will pass through about twice as many PNs in those cases.

If we subtract Equation (3) from Equation (4) and solve for N we obtain an expression for the reduction in N_{max} as a function of the number of routing nodes, N_r:

$$\Delta N_{max} = \frac{-kL_c}{T + L_c} \tag{6}$$

This means that for every two routing nodes that a message must pass through the maximum number of PNs it can pass through is reduced by one when $L_c = T$, which is very nearly the case for the values derived for T and L_c

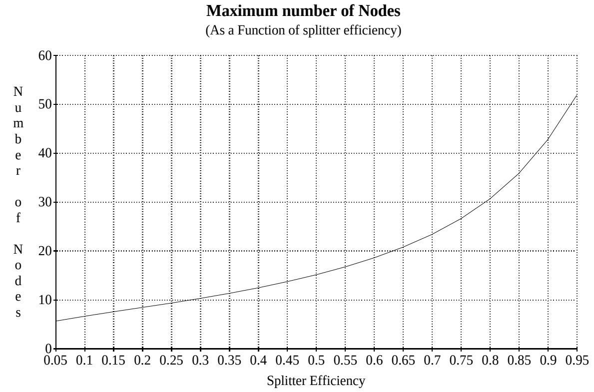

Figure 5. Maximum number of nodes as a function of splitter efficiency

above. For the conditions graphed in Figure 5 only six routing nodes can be accommodated. Therefore, a HORN of modest size can get by without any optical amplification, a more general, scalable system will require some intermediate amplification of the optical signals.

One undesirable affect ignored in the above consideration is that messages continue to circulate on the fiber rings until losses reduce their strengths to undetectable levels and that problem could be exacerbated by optical gain. Indiscriminate use of optical gain may result in an infinite gain condition, that is, the optical strength of a signal will grow exponentially if the gain for each pass around a ring exceeds the loss. It was therefore decided not to use gain on the local rings and to use gain only at the initial insertion point for remote transmissions. In this way the signal is amplified only once, not every time it passes a certain point on a ring.

Figure 6 shows signal strength versus the number of PNs in HORN for one and two full passes through the transmission path for a system in which the splitters at each node have a transmission efficiency of 90%. It shows again that losses don't reduce the signal below detectable levels for any path with less than 43 intermediate PNs and that the signal drops below detectable levels in less than two passes through any path with more than 22 intermediate PNs. This shows that signals can be made to die out due to normal losses within two passes while still ensuring sufficient signal strength to all nodes in the first pass, further supporting the decision to use gain elements only at the initial insertion point for remote transmissions.

2.3. Bit Error Rate

A second issue that should be addressed in conjunction with OPB is signal transmission bit error rate (BER) which is defined as the number of bits received in error divided by the number of bits transmitted. When the signal and noise sources are uncorrelated, that is they occur at random intervals relative to one another, then the BER is given by [17]:

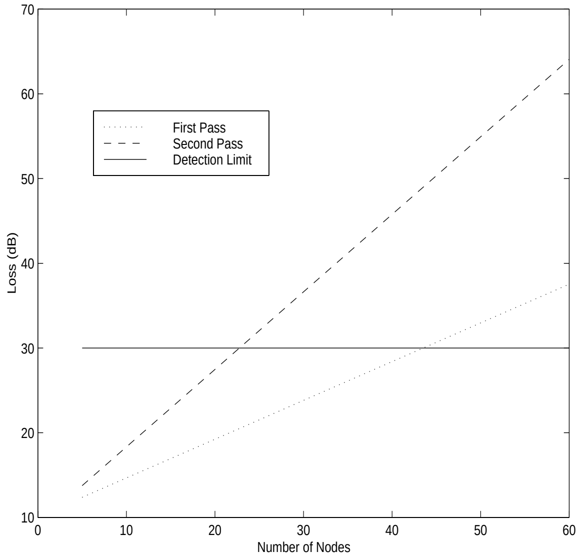

Figure 6. Losses in HORN as a function of the number of intermediate nodes. These results are valid for both local and remote transmissions if 10% of the signal is absorbed an each intermediate node. An intermediate node may be either a processing node or a routing node.

$$BER = \frac{\sqrt{2}exp(-\frac{SNR^2}{8})}{\sqrt{\pi}SNR} \tag{7}$$

In any electro-optic system random photon or electron fluctuations may introduce noise at any element in the system. Crosstalk at splitters and routers pose additional potential noise sources. When noise terms are uncorrelated, the SNR is given by:

$$SNR = \frac{I_{signal}}{\sqrt{\sum_j I_j^2}} \tag{8}$$

where the noise terms in the denominator, represented by I_j, are introduced by the source, detector, fiber and other intermediate elements.

The primary noise sources in the source and detector are shot noise, flicker noise and Johnson noise. The shot noise has an AC component and a DC component given by:

$$I^2_{shot(ac)} = 2qI_{signal}\Delta f \tag{9}$$

$$I^2_{shot(dc)} = 2qI_{DC}\Delta f \tag{10}$$

where q is the charge on an electron and Δf is the electrical bandwidth of the system. I_{signal} and I_{DC} are the average AC and DC currents in the system at the point where the noise is measured. The Johnson noise is given by:

$$I_t^2 = \frac{4kT}{R_{eq}} \Delta f \tag{11}$$

where k is Boltzmann's constant, T is temperature in degrees Kelvin and R_{eq} is the circuit resistance. Flicker noise is only present at low electronic frequencies and is negligible in HORN. The following values are reasonable and achievable for these equations subject to our need to minimize the impact of shot and Johnson noises on our network.

$$\begin{aligned} \Delta f &= 1MHz \\ I_{signal} &= 10\mu A \\ I_{DC} &= 10mA \\ R_{eq} &= 10k\Omega \end{aligned}$$

Crosstalk levels in AOTF's have been reported the literature at 15 dB below the signal levels for wavelengths spaced a few nanometers apart [11, 19]. The total contribution to the noise from the AOTF's will be a function of the number of routing nodes that a data transmission must pass through. In the worst case data will have to pass through N_r AOTF's. The noise contribution due to optical amplification will be on the order of 40 dB below the signal levels, which is the level reported for Willner and Hwang [20] for erbium doped fiber amplifiers (EDFA's) at wavelengths at and near 1550 nm. There will be, at most, one contribution to noise from the EDFA's in HORN since a transmission is amplified only when it is inserted onto the network for remote transmission. The summation in the denominator of equation (8) is therefore given by:

$$\begin{aligned} \sum_j I_j^2 &= I_{shot(ac)}^2 + I_{shot(dc)}^2 + I_t^2 \\ &\quad + k \times I_{AOTF}^2 + I_{EDFA}^2 \end{aligned} \tag{12}$$

where I_{AOTF} and I_{EDFA} are given by:

$$I_{AOTF} = 0.032 \times I_{signal} \tag{13}$$

$$I_{AOTF} = 0.0001 \times I_{signal} \tag{14}$$

Data circulating indefinitely around the different rings in HORN represents another potential source of crosstalk. In can be shown fairly easily that this is not a significant contribution to noise in HORN. First, the distances around any ring in HORN are so short that a message must make many loops around before it is out of phase to any significant extent with the basic signal. The choice of using splitters without gain elements therefore means that the signal has been attenuated by at least 20 dB before it experiences a measurable transmission delay relative to the signal in the first pass around a ring. Second, this transmission delayed component is highly correlated to the basic signal. While it may broaden the signal temporally, it doesn't qualify as noise.

It is therefore possible to achieve a BER better than 10^{-6} if the number of routing nodes that remote transmissions must pass through is limited to ten. This will permit a HORN with three levels to accommodate up to 2000 PNs. At present the crosstalk in the AOTF's is the limiting factor in this analysis.

3. Conclusion

While the application of optics to INs will theoretically vastly improve their performance, current limitations in optical technology severely limit the number of wavelengths that can be used simultaneously. Hierarchical topologies hold the promise of making optical INs practical in that they will permit hundreds, perhaps thousands of PNs to be interconnected using current technology. Several key issues facing hierarchical optical networks need to be resolved to establish their practicality and this paper addresses two of those issues, DCA and OPB, for a specific hierarchical optical IN, HORN, which was presented previously [15].

Four approaches to DCA were presented and discussed. The approach that requires full representation of the entire network at each node is unacceptable since the growth of the database and increase in communication traffic to maintain the currency of the databases makes that approach unscalable. It was also concluded that the number of tunable transmitters and receivers in the first three approaches was unacceptable. Multiple, tunable transmitters and receivers significantly increase the cost, and therefore adversely affect the scalability of HORN. While multiple transmitters and recievers may be beneficial under the last approach they are not required by the approach and can be fixed, and therefore much lower in cost. Improvements in delay and throughput due to reallocating channels dynamically have been found not to be as significant as initially hypothesized and there are indications that delays will actually increase if DCA is used [4]. These results coupled with the increased communication required maintain the databases in the first two approaches lead to the selection of the last DCA approach for use in HORN.

HORN was shown to be practical in terms of OPB and BER. It was shown that losses in HORN were low enough that optical gain elements are not necessary for local transmissions, i.e. transmissions that don't leave the lowest ring in HORN. Sufficient signal strength can easily be maintained in a system with up to about 43 intermediate PNs in the transmission path which is much less than limit of twenty imposed by current transmitter, detector and fiberoptic signal insertion/extraction technology. Use of optical gain at only the point that a signal is initially inserted onto the remote transmission rings permits signal strength to be maintained if much larger networks that follow the HORN topology while also preventing the open loop gain that would occur in some case when signals is amplified indiscriminantly every pass certain points on the rings since signals will circulate indefinitely when gain exceeds loss around a closed circuit. This also avoids the problem have having to remove signals actively from the rings. When splitter transmission efficiencies of 90% at are used each PN, losses were shown under this approach to signal strength by at least 10 dB for each circuit a signal makes around a ring. Noises in the sources, fibers, detectors and routing elements were shown to be very low. The contributions to the BER from those elements was much less than 10^{-50}. Similarly the contribution to the BER was shown to be very small due to a message continuously cycling through the system until losses reduce its strength below noise levels.

These results, coupled with the protocol analysis performed previously [12], show that HORN is practical and that the initial objectives of high throughput, low delay and high scalability can be achieved. These results also meet the demands of that variations in communication traffic may impose on an IN, avoid the need for actively removing signals, and ensure sufficient signal strength to all PNs in HORN even as elements and rings are added.

References

[1] M. Adda. A Scalable Multibus Configuration for Connecting Transputer Links. *IEEE Transactions on Parallel and Distributed Systems*, 8:245 – 263, March 1997.

[2] S. Dandamudi and D. Eager. Hierarchical Interconnection Networks for Multicomputer Systems. *IEEE Transactions on Computers*, 39:786 – 797, June 1990.

[3] P. Dowd, K. Bogineni, K. A. Aly, and J. A. Perreult. Hierarchical Scalable Photonic Architectures For High Performance Processor Interconnection. *IEEE Transactions on Computers*, 42:1105 – 1120, September 1993.

[4] P. Dowd, K. Bogineni, K. A. Aly, and J. A. Perreult. Addendum To Hierarchical Scalable Photonic Architectures For High Performance Processor Interconnection. *IEEE Transactions on Computers*, 43:864, July 1994.

[5] P. Dowd, J. A. Perreult, J. Chu, D. C. H. R. Minnich, D. Burns, F. Hady, Y. Chen, M. Dagenais, and D. Stone. LIGHTNING Network and System Architecture. *IEEE Journal of Lightwave Technology*, 14:1371 – 1387, June 1996.

[6] M. R. Feldman, S. C. Esener, C. C. Guest, and S. H. Lee. Comparison Between Optical and Electrical Interconnects Based on Power and Speed Considerations. *Applied Optics*, 27:1742 – 1751, 1988.

[7] J. W. Goodman, F. Leonberger, S.-Y. Kung, and R. A. Athale. Optical Interconnections for VLSI Systems. *Proceedings of the IEEE*, 72:850 – 866, July 1984.

[8] M. Hamdi and S. W. Song. Embeding Hierarchical Hypercube Networks Into The Hypercube. *IEEE Parallel and Distributed Systems*, 8:897 – 901, September 1997.

[9] P. A. Humblet and W. M. Hamdy. Crosstalk Analysis and Filter Optimization of Single and Double Cavity Fabry Perot Filters. *IEEE Journal on Selected Areas in Communications*, 8:1095 – 1107, June 1990.

[10] E. Iannone and R. Sabella. Optical Path Technologies: A Comparison Among Different Cross-Connect Architectures. *IEEE Journal of Lightwave Technology*, 14:2184 – 2196, October 1996.

[11] J. L. Jackel, M. S. Goodman, J. E. Baran, W. J. Tomlinson, G.-K. Chang, M. Z. Iqbal, G. H. Song, K. Bala, C. A. Brackett, D. A. Smith, R. S. Chakravarthy, R. H. Hobbs, D. J. Fritz, R. W. Ade, and K. M. Kissa. Acousto-Optic Tunable Filters (AOTF's) for Multiwavelength Optical Cross-Connects: Crosstalk Considerations. *IEEE Journal of Lightwave Technology*, 14:1056 – 1066, June 1996.

[12] T. S. Jones and A. Louri. Analysis of Collision-Free, Single Hop Media Access Protocols For A Scalable Optical Interconnection Network. *Submitted to IEEE Journal of Lightwave Technology*, 1998.

[13] L. G. Kazovsky and P. T. Poggiolini. STARNET: A Multi-Gigabit-Per-Second Optical LAN Utilizing A Passive WDM Star. *IEEE Journal of Lightwave Technology*, 11:1009 – 1026, 1993.

[14] G. Keiser. *Optical Fiber Communication*. McGraw-Hill Book Company, New York, 1991.

[15] A. Louri and R. Gupta. Hierarchical Optical Interconnection Network (HORN): Scalable Interconnection Network for Multiprocessors and Multicomputers. *Applied Optics*, 36:430 – 442, January 1997.

[16] M. A. Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri. MAC Protocols and Fairness Control in WDM Multirings with tunable transmitters and fixed receivers. *IEEE Journal of Lightwave Technology*, 14:1230 – 1244, June 1996.

[17] C. R. Pollock. *Fundamentals of Optoelectronics*. Irwin, Chicago, 1995.

[18] K. Sivalingam and P. Dowd. A Multilevel WDM Access Protocol For An Optically Interconnected Multiprocessor System. *IEEE Journal of Lightwave Technology*, 13:2152 – 2167, November 1995.

[19] D. A. Smith, R. S. Chakravarthy, Z. Bao, J. E. Baran, J. L. Jackel, A. d'Alessandro, D. J. Fritz, S. H. Huang, X. Y. Zou, S. Hwang, A. E. Willner, and K. D. Li. Evolution of the Acousto-Optic Wavelength Routing Switch. *IEEE Journal of Lightwave Technology*, 14:1005 – 1018, June 1996.

[20] A. E. Willner and S.-M. Hwang. Transmission of Many WDM Channels Through a Cascade of EDFA's in Long-Distance Links and Ring Networks. *IEEE Journal of Lightwave Technology*, 13:802–816, May 1995.