Pushing Away the Communication bottleneck with Optical Interconnects in Symmetric Multiprocessors

Wissam Hlayhel^{*a}, Jacques H. Collet^b, Christine Rochange^a, and Daniel Litaize^a

^aInstitut de Recherche en Informatique
118 route de Narbonne, Toulouse Cedex 4, 31062 France.

^bLaboratoire d'Analyse et d'Architecture des Systèmes du CNRS
7 avenue du Colonel Roche, Toulouse Cedex 4, 31077 France.

ABSTRACT

We analyze the bandwidth needed for transmitting the addresses in future symmetric multiprocessor machines (SMP), constructed around a shared bus due to the critical obligation to preserve the coherence of the memory hierarchy. We show that an address-transaction bandwidth as high as several hundreds of Gbit/s will be necessary not to slow down the execution of most applications in large SMP's. This communication bandwidth seems incompatible with the operation constraints of shared electrical busses, making necessary the search for other implementations of the address transmission network. We consider the introduction of optical interconnects (OI) in this context. We review several solutions, in the ascending order of complexity of the optical subsystems as one critical issue concerns the degree of sophistication of the optical solutions and their cost. We first consider simple point to point OI's for a SMP chipset. The interest for OI's comes from the low energy consumption and from the possibility, in the future, to integrate several thousands of optical input/outputs per electronic chip. Then we consider the implementation of an optical bus that is a multipoint optical line involving more optical functionality. We discuss the possibility of multiple accesses to the bus, and the constraints related to the necessity to maintain the coherence of caches.

Keywords: Symmetric Multiprocessors, Shared busses, Bandwidth, Optical Interconnects, cache coherence

1. INTRODUCTION

The logical way to reach a performance level not accessible with a mono-processor consists of connecting several processors through an interconnection network. One must distinguish distributed multicomputer systems (DMS) from tightly bound multiprocessors (TBM). A DMS may include from a few tens to several hundreds of computers separated by an intercomputer distance ranging from a few meters to a few hundreds of meters. They are best suited for transmitting steady data streams between nodes (video, files of a few Mbytes, etc.). DMS's are the field of numerous investigations partly motivated by the preoccupation of developing more modular, low cost hardware that would simplify maintenance and compatibility issues for manufacturers. However, they are suitable for some, but not all applications because the latency of internode communications becomes extremely long (with respect to the processor cycle, 1ns today) when the inter-computer distance attains a few meters (1meter=5ns). Therefore, a DMS will execute much slower many applications (especially those requiring numerous internode exchanges) than a TMB enclosed in a single cabinet.

We focus in this work on the communication challenges in TBM's and on the possible role of optical interconnects (OI's) in this context. The processor distance in these machines ranges from a few centimeters to a few tens of centimeters. The latency is the key parameter of the memory-processor interactions because the processors exchange very short bursts of information (usually at least one word, i.e., four bytes, more often a cache block at a time, i.e., 32 or 64 bytes). They never establish a steady communication data stream with the memory. Although hundreds of network topologies have been proposed, the number of commercial implementations is handful and mostly reduces to shared busses, rings, meshes, tori and central switches. Obviously, all topologies are of equal conceptual interest. However, symmetric multiprocessors (SMP) represent a non-negligible fraction of the market, contrarily to other topologies, which are practically very marginal. SMP's are build with a shared bus, which grants a uniform access time to the shared memory, provides a single addressing space

* Correspondence: Email: hlayhel@irit.fr; collet@laas.fr; litaize@irit.fr;

and enables the programmer to get transparent access to computational resources. In SMP's, it is possible to distinguish three contributions to the memory access latency, namely:

- 1) The *intrinsic memory latency*, typically today of the order of 40-60 ns.
- 2) The *bus contention latency*, which critically depends on network saturation. The bus bandwidth directly controls the contention latency depending on the number of connected processors.
- 3) The *coherence latency*. Maintaining the coherence of the caches is absolutely mandatory in tightly bound machines and requires broadcasting (or multicasting) coherence messages through the communication network. This coherence preservation contributes to slowing down the memory access. Snooping protocols in bus-based machines rest on the property that the shared bus represents a synchronisation point watched by all processors, enabling simultaneous update or invalidation of duplicated copies of data in caches¹.

Figure 1. Example of a modern SMP architecture with 4 processors, 2 address busses (Adr bus 1, Adr bus 2) and the associated processor caches (C1 for Adr bus 1, C2 for Adr bus 2).

The paradox of SMP's is that the shared bus is simultaneously an enormous advantage and a terrible weakness. It is particularly advantageous as it enables implementing snooping protocols, which are much simpler and faster than directory-based protocols that have to be used in distributed networks². It is also a terrible weakness because the bus serialises the communications with the memory and therefore represents a transmission bottleneck, which increases the memory access latency. Modern machines mitigate this issue by separating the transfers of addresses and data as will be shown below. However, the bandwidth of the electronic address bus represents the most problematic communication bottleneck of current and future SMP's. Although the transfer of data requires much more bandwidth, it is not a real issue as each processor can be connected to the crossbar switch by a private link (see figure 1), the bandwidth of which can be increased with no coherence constraints.

The rest of the work is organised as follows: 1) We describe the current state of SMP's, particularly the communication performance; 2) Using a multiprocessor simulator, we study the address bandwidth requirements of the next generation of SMP's. Considering the performance of future superscalar processors, this study shows that an address-transaction bandwidth of 80-100 Gbit/s will be necessary not to slow down the execution of applications in future large SMP's (including for instance 64 processors operating at 1 GHz). Reaching such a high bandwidth is particularly difficult using electric solutions because of the operation limitations of shared electric busses. It seems necessary to search for other implementations of the address transmission network in future SMP's; 3) We consider the introduction of OI's in this context. We review several solutions in the ascending order of complexity of the optical subsystems, as one critical issue concerns the degree of sophistication, the feasibility and the cost of the optical solutions. We first consider the sole introduction of point to point OI's in the framework of a SMP chipset. The interest for OI's comes from their low energy consumption and small size enabling the implementation of several thousands of optical input/outputs (IO's) per electronic chip as needed by a multiprocessor chipset. Then we consider the introduction of an optical U-bus, which operates in the synchronous mode. The U-bus operates with no access arbitration and enables simultaneous transfers of different

transactions to the bus. However, the price to pay is that the coherence of caches cannot be preserved by means of the usual snooping protocols (MESI protocol). New cache coherence protocols are under investigation.

2. CURRENT SMP ARCHITECTURES

The fast increase of the processor speed has induced a dramatic growth of the bandwidth needs in recent SMP's as can be shown from the evolution of shared-bus machines. For instance, the *XDBus* from Sun³ is a 64-bit multiplexed address-data bus running at 40 MHz (10 slots) which involves a split-transaction protocol. Each bus transaction requires 2 bus cycles for an address transfer and 9 bus cycles for a cache block transfer. The raw bandwidth is 320 MB/s, but the effective bandwidth is only $\frac{3}{4}$ of this value. In the SGI Power Challenge⁴, the split-transaction protocol is enhanced to take into account the presence of two busses clocked at 47.6 MHz (13 slots), one of 40 bits dedicated to the address transfer and another one of 256 bits used for the transfer of cache blocks. The two busses are synchronized and enable to transfer the address and the data in 5 bus cycles. The busses can support up to 8 outstanding read requests. The raw bandwidth is 1.5 GB/s but the effective bandwidth reduces to $\frac{4}{5}$ of this value. The *Sun Enterprise 6000*⁵ is also a non-multiplexed split-bus running at 83.5 MHz with 256 bits for the data and 41 bits for the address. This high operation frequency (for a shared bus) is due to a special backplane design (so-called "Centerplane" technology) where cards are connected by both sides (8 slots per side). The raw bandwidth is 2,6 GB/s and the number of outstanding read requests that can be in progress at the same time is 112 and up to 7 from each board, each processing board containing two processors). In recent SMP designs, such as the *G30* machine from IBM or the 10000 from SUN⁵, the data network is a crossbar, but the address network remains a shared bus used for maintaining coherency. Above a certain number of processing elements, multiple address busses are needed (see Figure 1). For example, two XDbusses were implemented in the *SparcCenter 2000* to attain a maximum number of 30 processors, four busses in the *Enterprise 10000* to attain 64 processors. With a smart design of the crossbar switch, the direct transfer of blocks from cache to cache becomes possible without updating the main memory. The most critical problem at the moment is the address/snoop bandwidth discussed in the next section.

Figure 2. Execution time versus bus bandwidth for different applications of the SPLASH suite. The execution times of the different applications are normalized to 1 when considering an infinite bus bandwidth.

3. ADDRESS BUS TRAFFIC AND LIMITATIONS OF ELECTRIC BUSSES

We studied the address bus traffic using an instruction-driven simulator that enables tracing the bus operation during the execution of standard applications. The main features of superscalar processors (i.e., multiple instruction issue, dynamic scheduling, non-blocking reads, register renaming and speculative execution) are taken into account by the simulator and

can be modified by means of adjustable parameters. Reference 6 describes in details the simulator and previous simulation results, showing in particular the address traffic traces, which reveal the high irregularity of communications. To characterize the impact of the address bus bandwidth on the execution time of the applications, we assumed that the network and the memory are contention-free so that the network transmits all the requests at each cycle.

The main practical conclusions are reported in Figure 2, which displays the execution time of different applications (Radix, Ocean, LU, FFT) of the SPLASH suite ⁷ as a function of the bus bandwidth. We assumed an intrinsic memory access time of 75 processor cycles, that is equivalent in practice to considering processors running around 1-1.2 GHz. For clarity, all execution times have been normalized to 1 when considering an infinite bus bandwidth. Figure 2 shows that the execution time of most applications very weakly increases when the address bus can transmit from 0.5 to 1 request per processor cycle (RPC). The worst case is that of the *radix* application, with an execution time approximately multiplied by 2 when the bus bandwidth is 0.5 RPC. Now, each address transaction generally requires 4-5 bus cycles due to the operations of access arbitration, transmission and preservation of the coherence (see section 2). Thus, for processors operating at F=1 GHz, the transmission of 0.5 RPC requires a bus operating 4 to 5 times faster (i.e., around 2- 2.5 GHz). The top bus bandwidth for 40-bit addressing will be in the range of 80-100 Gb/s. To avoid any confusion, it must be stressed that this bandwidth is not a point to point bandwidth, but that of a multipoint line, with the additional requirement that the propagation time must be shorter than one bus cycle (1 ns in our example) to be consistent with the bus operation principle. These constraints seem far beyond the capabilities of electric shared busses today based on the Gunning Transceiver Logic (GTL) technology. They operate as a transmission line ^{8 9}. Two points are important:

- The propagation time of logical signals along the electric bus is $\tau = \sqrt{LC}$, with L and C being the linear inductance (typ. 5-10 nH/cm) and the linear capacity of the bus, respectively. The capacity is mostly due to the transceivers that connect the different processors, each transceiver contributing typically for $C_t=20$ pF. The propagation time when connecting N processors over the length l is $\tau = \sqrt{LNC_t/l}$ (Eq. 1). For instance, τ equals 25 ns/m when one processor is connected every second cm.
- The bus operation frequency F_B must be low enough to ensure that the bus has time to attain an equipotential state (despite propagation effects) during the bus cycle.

With these constraints, even assuming a perfect adaptation of the termination load and no line reflection, the operation frequency of the bus cannot exceed $F_{MAX} = 1/(l\tau)$, or using Eq. 1, $F_{MAX} = 1/\sqrt{lLNC_t}$ (Eq. 2). This formula leads to very profound limitations. For instance, a 1.28m-long bus connecting 64 processors (with $C_t=20$ pF) could not operate at a frequency higher than 30 MHz. The address bandwidth, assuming 41-bit address transmission as in the SUN Enterprise 6000, would be of the order of 1.2 Gb/s, i.e., about two orders of magnitude lower than that deduced from the simulations! The only “palliative” solution consists of duplicating the number of address busses (for instance considering 10 or 20 bus) to reduce the operation frequency for each one. With N_B busses transmitting with a parallelism N_W (typically 40), the accessible bandwidth B reads: $B=N_B N_W F_{MAX}$. Using Eq. 2, The maximum bus length to transmit the bandwidth B with N_B busses is:

$$l = \frac{N_B^2 N_W^2}{LNC_t B^2} \quad (3)$$

Let us consider a numerical example with $B=100$ Gb/s (result of simulations), $N_W=40$ (bus width), $N=64$ (nb of processors), $L = 5$ nH/cm, $C_t = 20$ pF/transceiver. We deduce $l \approx 2.4 \times 10^{-2} N_B^2$ with l in cm. Thus, with $N_B=10$ busses, the top acceptable bus length would be 2.4 cm to connect 64 processors! The corresponding interprocessor board distance shorter than 1 mm seems unrealistic. Moreover, the system will be complicate, as each processor would have to be connected to 10 caches and to 10 arbiters (one per bus). To go around this issue, one may possibly group several processors per daughter board. For instance, using quadriprocessor cards, the number of connected cards reduces to $N=16$. The bus length $l=10$ cm becomes possible, with an interboard distance of 6 mm. Of course, neither the large number of busses, nor their necessary short length, nor some additional energy dissipation issues are insuperable obstacles that will block in the future the feasibility of fully electronic SMP's, but a technological break seems necessary. We discuss in section 4 several strategies for introducing OI's in SMP's.

4. IMPLEMENTATION OF OPTICAL COMMUNICATIONS IN SMP'S

There are many technological arguments to suggest the introduction of OI's^{10 11}. Perhaps, the two following ones are mostly important:

- Transmitting a few Gb/s with OI's is much simpler than with electric lines. There are no line adaptation problems, no capacity effects to slow down the propagation velocity. Remember that the propagation time equals 5 ns/m in the optical fibers, that is much less than the value of 25 ns/m deduced previously in the GTL bus with one transceivers connected every second cm, see Eq. 1. Thus, an optical bus (that is of course an optical system much more complicated than a simple point to point optical link) could *a priori* operate at higher frequency than an electric one.
- The switch energy per bit is lower for OI's than for electric lines. The energy dissipated for transmitting one bit in an electric line is $E = C_{IO}V^2$, where V is the chip operation voltage and C_{IO} the IO capacity, mainly due to packaging. The study of the evolution of VLSI's shows that C_{IO} is of the order of 10 pF and does not evolve with time, despite the reduction of the lithographic features and the size of transistors in the chips. Assuming 2V for the voltage operation, the switch energy is 40 pJ/bit. Let us consider the switch energy of an optoelectronic interface. To day, VSCSEL's with a threshold current of 0.1 mA have been demonstrated. They can operate with a modulation current of 0.2 mA. At 1 GHz, the emission energy per bit is of the order of 0.2 pJ/bit (i.e.: 0.2 mA * 2volt * 500ps). The estimation of the energy involved in the reception stage is more complicated and depends on the detector capacity C_{PD} (that is proportional to the detector area). C_{PD} is typically of the order of 0.5 fF/ μm^2 . There is a necessary tradeoff between *making the detector as small as possible* (to run fast and to minimize the consumption) and *maintaining an acceptable coupling efficiency* of the diode with the external incoming beam that requires a minimum detector surface. Assuming for instance a detector area of 20x20 μm^2 and a voltage operation of 2V, the detector switch energy is 0.8 pJ/bit (again $E = C_{IO}V^2$). The total switch energy is of the order of 1pJ/bit, that is to say much smaller than that deduced for the electric interconnect. One may expect an important reduction of the energy consumption of interchip communications. We stress that reducing the energy dissipated in interchip communications is crucial for the evolution of VLSI circuits as the energy is more and more dissipated by the communications rather than by the logical processing¹².

We review several solutions, in the ascending order of complexity of the optical subsystems as one critical issue for the viability of optical solutions concerns their degree of sophistication and their cost.

Figure 3. New SMP implementation including a central multiprocessor chipset with optical links connecting the chipset to the processors and the memory

4.1. OI's for a SMP chipset

In fact, most of the limitations of the electric bus arise from the sole fact that it is a transmission line, which slowly reaches the equipotential state due to the capacity of the connected transceivers. Thus, integrating the whole bus (or possibly different busses, as shown in Figure 1) in a single VLSI would dramatically diminish the electric constraints, particularly those resulting from the "equipotentiality" condition of the bus. Note that integrating the bus does not necessarily imply maintaining the bus architecture in the new chip (particularly the presence of long parallel electric lines). One may consider constructing a logical circuit that mimics the bus operation, mainly the mechanism of arbitration, the serialization of the address transactions with the memory, and the preservation of the cache coherence. Replacing the address bus by a central

switch is nothing but the continuation of an idea already put into practice in the monoprocessor architecture with the chipset integration. However, the generalization to multiprocessors with the definition of a multiprocessor chipset (MPC) is not simple for several technological reasons. The most important obstacle is the chip connectivity. Connecting N processors to the MPC would require at least $64 \cdot N$ pins (assuming 64-bit wide transmissions), plus $64 \cdot M$ extra connections to M memory chips, plus some more pins for the power supply, ground., etc. With $N=32$ processors, $M=4$ memory chips, the number of pins would be around $N_p=2300$. This estimation does not include the data transfer network (i.e., the private links to the crossbar) so that integrating both the data and the address network would require of the order of 5000 pins. The mechanical feasibility of such a high number of pins seems so far not demonstrated.

OI's might contribute to better solve the problem. The vertical emission of VCSEL's enables integrating optical IO's anywhere on a VLSI chip, not only on the perimeter. Arrays of 8×8 emitters have been demonstrated using the hybrid integration of III-V VCSEL's on top of CMOS $0.25 \mu\text{m}$ drivers¹³. Similarly, arrays of Si-diode detectors with their trans-impedance amplifiers have been also integrated in CMOS circuits. The inter-emitter (or inter-receiver) distance of the order of $400 \mu\text{m}$ enables considering a density of about 600 optical IO/ cm^2 . Therefore, OI's represent a technological breakthrough that could permit to reach the IO densities in the range of several thousands necessary for a MPC. This will be also possible because the low switch energy of OI's (0.2 mW for a VCSEL operating at 1 GHz with a current of 0.2 mA) would dramatically contribute to reducing the energy consumption of the chip. There are several solutions to connect the emitters and the receivers of different chips. One may consider free space interconnects¹⁴, or guided transmissions through an optical backplane¹⁵ or ribbons of fibers. The different solutions seem viable, but the simplicity of implementation, the mechanical stability (that is a critical feature of optical systems) and above all the cost will be the final decisive arguments.

Figure 4: Example of an Optical folded bus (U-bus) with 5 nodes and 2 memory chips. A node can be a mono or a multiprocessor card. Packets are added in the *LOAD zone* and move to the right towards memory chips.

4.2. Optical bus

The MPC approach suggested in the previous paragraph only requires point to point optical interconnects and is based on the underlying idea of extending the bus bandwidth by electronic integration without questioning the basic operation principle of the system, in particular the serialization of the address transactions to the memory. Another alternative to go around the address bandwidth bottleneck in future SMP's might consist of relaxing the bus access arbitration. Two options are possible:

- Considering this approach in the framework of the MPC. It is possible to consider the integration of a true optical bus with the tools of integrated optics, i.e., optical guides, Y switches, etc... It is also possible to consider a pure electric solution with a new design of a MPC architecture that would enable simultaneous multiple accesses.
- Implementing of a real optical U-bus as shown in Figure 4^{16 17}. The bus operation is described below. Naturally, the construction of an optical bus (on-chip or off-chip) implies an increase of the involvement of optical technologies with respect to the MPC approach described in section 4.1. The bus is a multipoint optical line. In particular, it requires that each node adds and drops optical data with concomitant energy loss and energy balance issues.

The U-bus may operate in the synchronous and asynchronous modes. The synchronous operation is based on two phases, namely:

- The LOAD phase. Each processor inserts (if needed) one optical pulse in the bus *LOAD zone* (see Fig. 4). Although the pulse duration depends on the optical emitter technology, duration as short as a few tens of picoseconds is easily achieved today with mode-locked lasers or fibers. The only geometric restriction to avoid access arbitration is that the internode distance L must be longer than the optical pulse (i.e., $L > cW_p$ where W_p is the duration of the pulse). In that case, there is no data overlapping from adjacent processors during the load phase. Considering a bus as large as the address size makes possible transmitting one full address per cycle.
- The propagation phase. It follows the LOAD phase and lasts until the address inserted by the first node bypasses the last one (respectively nodes 1 and 5 in figure 4). This time T_p is necessary to free completely the *LOAD zone* from data. T_p depends on the length of the *LOAD zone*, and can be as short as $N.W_p$, where N is the number of connected processors. Thus, an optical bus may potentially transmit one address to the memory every W_p seconds. For instance with 10-ps pulses and $N=64$ processors, the bus access latency could be as short as 640 ps and the bandwidth as high as 100 GAddress/s. This extremely huge bandwidth is too large for existing memories, but it demonstrates the optical bus potentiality. In practice, the solution to match the bus bandwidth with that of the memory consists of moving away the nodes to slow down the bus traffic (also increasing the bus access latency). For instance, with one node every 20 cm, the U-bus bandwidth reduces to 1 GAddress/s.

Bus operation up to a few GHz (or higher) becomes possible (in replacement of electric bus operating around a few tens of MHz) because the transmission of optical pulses in guides is not penalized by capacity effects and critical load adaptations encountered for electrical transmissions in a multi-point line. Moreover, the parallel transmission through optical lines is almost skew-free in the GHz domain for transmission over a few tens of meters. This simplifies data sampling in case of parallel transmissions. As a result, the SMP architecture (i.e., the processor, the bus and the memory) would become more scalable.

Dropping the access arbitration to the U-bus simplifies the implementation and extends the bandwidth, but unfortunately there are also several bad consequences. The price to pay is very high in term of system operation, due to impossibility to maintain the coherence of caches using the snooping protocols of standard busses². It must be stressed that maintaining the coherence of caches is absolutely mandatory to preserve the machine performance. Dropping the coherence would have two dramatic consequences:

- A very important increase of the bus traffic as each processor would have to fetch the data operands in the memory (caching shared data would not be allowed).
- A dramatic increase of the data access time because of the obligation to read in the memory. In the presence of the full hierarchy, the average data access time is of the order of 2-3 processor cycles, and would be close to 100 cycles without caches.

These simple considerations demonstrate that changing the SMP architecture might be meaningless if the coherence of caches cannot be preserved. Of course, maintaining the coherence is always possible with the implementation of directory-based protocols, but they are much slower than the snooping solutions due to the increase of the coherence traffic. At the moment, it is not demonstrated that the advantages gained by dropping the arbitration access in an U-bus will offset the consequences related to the implementation of directory-based protocols. A new snooping solution might be possible as the U-bus serializes the accesses to the memory so that all processors can simultaneously watch the memory transactions at point S in figure 4. Nevertheless, the issue comes from the fact that the atomicity of transactions is not preserved that makes more complicated the actions to maintain the coherence of caches. Adapted coherence protocols are under investigation.

5. CONCLUSION

We showed that the increase of the processor power leads to dramatic bandwidth needs for future SMP's, especially concerning the address transmission that is complicated by the obligation to maintain the coherence of caches. Neither the increase of the bus number, nor their necessary small length, nor some possible energy dissipation issues are insuperable obstacles that will block in the future the feasibility of fully electronic SMP's, but a technological break seems necessary. The integration of OI's can help in designing new SMP's. We considered two possible solutions.

- The MPC approach requires only point to point optical interconnects and is based on two underlying ideas: 1) Including the minimum of optical technology that enables to improve significantly the communication system, due to the economic risk factor; 2) Extending the bus bandwidth by electronic integration in a single VLSI without questioning the

basic operation principle of the system, in particular the serialization of the address transactions to the memory. The interest of OI's comes from the low energy consumption and the small size enabling the integration of several thousands of optical input/outputs per electronic chip needed by the MPC.

- The optical bus approach based on the construction of a multipoint optical line. This alternative requires more sophisticated optical components, which often are expensive and not very integrated at the moment. Optical solutions with ribbons of Y switches are not available. The only possible solution seems to use an optical backplane with diffractive gratings for coupling the different processors to the optical bus. It is clear that this technology is at the moment expensive and in a preliminary development stage. The most critical question for the viability of this approach is the preservation of the coherence of caches that is mandatory for the efficient operation of an SMP.

ACKNOWLEDGEMENTS

This work was conducted in the framework of the *Optical Interconnects for Integrated Circuits* project (OIIC) supported by the *MicroElectronic Advanced Research Initiative* (Mel ARI) of the European Communities.

REFERENCES

1. Mark Papamarcos, and Janak Patel, "A Low Overhead Coherence Solution for Multiprocessors with Private Cache Memories," in proceedings of the 11th Annual International Symposium on Computer Architecture, pp. 348-354, June 1984.
2. J.L. Hennesy et D.A. Patterson, "Computer Architecture, a Quantitative Approach," Morgan Kauffmann Publishers, second edition, 1996
3. Pradeep Sindhu, Jean-Marc Frailong, Jean Gastinel, Michel Cekleov, Leo Yuan, Bill Gunning and Don Curry, "XDBus: A High-Performance, consistent, Packet-Switched VLSI Bus," in proceedings of COMCON, Spring 93.
4. Mike Galles and Eric Williams, "Performance Optimisations, Implementation and Verification of the SGI Challenge Multiprocessor," Proceedings of 27th Annual Hawaiï International Conference on Systems Sciences, January 1993.
5. Technical Report: www.sun.com/datacenter/products
6. W. Hlayel, D. Litaize, L. Fesquet, and J.H. Collet., "Optical versus electronic bus for address-transactions in future SMP architectures," Proc. International Conference on Parallel Architectures and Compilation Techniques ; PACT 1998, Paris october 1998, IEEE Computer Society, p 22-29
7. S.C. Woo et al., "The SPLASH-2 Programs : Characterization and Methodological Consideration," Proc. International Conference on Computer Architecture (ISCA), Santa Margherita Ligure, Italy, 22-24 June 1995, pp.24-36.
8. Texas Instrument , "Fast GTL Backplanes with the GTL1655," Application report: SCBA015;january 1999
9. H. Johnson et M. Graham, "High speed digital design," Prentice Hall PRT, 1993, ISBN: 0-13-395724-1
10. D.A.B. Miller, "Physical reasons for optical interconnection," International Journal of Optoelectronics, Vol 11, No 3, p155-168, (1997)
11. G. Yayla, P. Marchand, and S. Esener, "Energy and speed analysis of digital electrical and free-spaceoptical interconnections," in Optical Interconnections and Parallel Processing: The Interface: Edited by A. Ferreira and P. Berthome eds., Kluwer, chap. 3, (1997)
12. See energy dissipation calculations in the ASIC data books of the IBM technical library at the URL: <http://www.chips.ibm.com/techlib/products/asics/databooks.html>
13. See the following sites:
MicroEletronic Advanced Research Initiative of the European Communities.: <http://www.cordis.lu/esprit/src/melari.htm>.
Optical Interconnected Intergrated Circuits: <http://www.elis.rug.ac.be/~jvc/oiic/sysdemo.htm>)
14. I. Redmond, and E. Schenfeld, "Experimental results of a 64 channel, free space optical interconnection network for massively parallel processing", Institute of Physics series number 139, proc. Optical Computing conference, pp 149-152 (1994)
15. H. Scott Hinton, and T.H. Szymanski, "Intelligent Optical Backplanes," IEEE proceeding of the 2nd international conference on Massively Parallel Processing using Optical Interconnects (MPPOI'95); pp133-143, (1995)
16. C. Qiao, and R. Melhem, "Time Division Optical Communications in Multiprocessor Arrays," ,IEEE Trans. on Computers vol 42, 577-590 (1993)
17. A. Guo, R. Melhem, R. Hall, D. chiarulli, ans S. Levitan, "Array processors with pipelined Optical busses," , J. Parallel Distributed Computing vol 12, 269-282, (1991)