



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*L'effet du degré de réseau d'interconnexion
multi-étages sur ses performances: le cas des réseaux
Delta et réseaux Delta surdimensionnés*

Ahmad Chadi ALJUNDI — Jean-Luc DEKEYSER

N° 4880

1^{er} juillet 2003

THÈME 1



R
apport
de recherche



L'effet du degré de réseau d'interconnexion multi-étages sur ses performances: le cas des réseaux Delta et réseaux Delta surdimensionnés

Ahmad Chadi ALJUNDI* , Jean-Luc DEKEYSER*

Thème 1 — Réseaux et systèmes
Projet Dart

Rapport de recherche n° 4880 — 1^{er} juillet 2003 — 16 pages

Résumé : Lors de la construction d'un ordinateur parallèle, le réseau d'interconnexion ainsi que ses performances sont des facteurs cruciaux à étudier. Aujourd'hui, la technologie supporte la construction de crossbars de taille limitée (jusqu'à 128). Pour des machines avec un plus grands nombre de points de connexion, ces crossbars peuvent être utilisés comme éléments de commutation de réseaux d'interconnexion multi-étages. En fait, un réseau d'interconnexion est défini par sa topologie, son algorithme d'acheminement, sa stratégie de commutation et son mécanisme de contrôle de flux. Un des facteurs définis par la topologie est le *degré*. Le degré d'un réseau d'interconnexion multi-étages est la taille des commutateurs qui le composent. Dans ce papier, nous étudions l'effet de la taille du commutateur sur les performances de deux classes de réseaux d'interconnexion multi-étages: les fameux réseaux Delta et une sous-classe de cette famille qui s'appelle les réseaux Delta surdimensionnés. Cette étude initie une réflexion sur l'utilisation des réseaux d'interconnexion multi-étages dans les machines SMP actuelles.

Mots-clés : Réseau d'interconnexion, réseau Delta, réseau Delta surdimensionné, évaluation de performances

* Université des Sciences et Technologies de Lille, Laboratoire d'Informatique Fondamentale de Lille, 59650 Villeneuve d'Ascq, France

The Effect of the Degree of Multistage Interconnection Networks on their Performance: the Case of Delta and Over-sized Delta Networks

Abstract: Interconnection network performance is a key factor when constructing parallel computers. Today's technological progress makes it possible to build and use crossbars of size up to 128. Such crossbars can be used as switching elements in parallel architectures such as multistage interconnection networks. In fact, A multistage interconnection network is usually defined by its topology, routing algorithm, switching strategy, and flow control mechanism. One of the factors defining the topology of a multistage interconnection network is its *degree*. The degree of a multistage interconnection network is the size of the switching element of which it is composed. In this paper we are interested in studying the effect of the size of the switching element on the performance of two classes of multistage interconnection networks: the famous Delta networks and a subclass of this family called the over-sized Delta networks. This study is to be used in future work in order to evaluate the use of multistage interconnection networks as an intercommunication medium in today's Symmetric Multiprocessors.

Key-words: Interconnection network, delta network, over-sized delta network, performance evaluation

1 Introduction

As electronic components reach their physical speed limits, using parallelism seems to be the only known practical solution to face today's scientific applications augmenting need for computation speed. In a parallel computer the intercommunication between processors and their communication with memories is a key factor on which the performance of the overall system depends. The main interest in an intercommunication system is that it has the capacity to route many communication tasks concurrently. A conflict occurs when more than one message try to access a same communication resource. Three types of conflicts[KD97] exist while accessing the shared memory in a multiprocessor system : network conflicts, bank busy conflicts, and simultaneous bank conflicts.

While a bus allows a very limited level of parallelism, a crossbar[WB72], which provides a full connection between all the nodes of the system is very complex, expensive and hard to be controlled. Interconnection networks[SY94], and more precisely Multistage Interconnection Networks[Sie77] (MINs) are an interesting solution for intercommunication systems in parallel computers. They provide acceptable degree of parallelism with a complexity well inferior to this of a crossbar. MINs are usually used in multi-processor and multi-computer parallel machines as an intercommunication medium between processing elements and memory modules. Many MINs belonging to the famous Delta MINs family were studied and effectively used to build parallel computers. Delta MINs form a sub-group from a bigger MINs family called banyan MINs of which networks are characterized by the existence of one and only one path between each source and destination. Non-banyan MINs are, in general, more expensive than banyan networks and more complex to control. Still, they often are fault tolerant and capable to apply rerouting strategies used to bypass problems found by messages such as a conflict or a faulty link or switch. Kruskal and Snir[KS83] studied *augmentation* techniques that might be applied on banyan networks in order to improve their performance without much loss in simplicity. Augmentation can be defined as adding links and/or switches to the basic configuration of the network.

A MIN is usually defined by[CSG99], among others, its topology, routing algorithm, switching strategy, and flow control mechanism. One of the factors defining the topology of a MIN is its *degree*. The degree of a MIN is the size of the switching element (SE) of which it is composed. In this paper we study the effect of the MIN's degree on its performance. Networks of degrees 2, 4, and 8 were studied by Cheemalavagu and Malek in [CM82]. On their paper, they were limited to 8 degree MINs because of the space and time needed for the simulation. Furthermore, they used networks with and without buffers. In this paper we investigate MINs of degrees up to 64, all of them unbuffered. In order to establish this study, different degrees Delta and over-sized Delta MINs[ADS03] are tested. The test methodology is a simulation one based on a *Universal Performance Factor (UPF)* explained in earlier publications[ADKS02, ADKS03, ADS03].

SMP (Symmetric MultiProcessing) machines are interesting architectures used today to build parallel computers. They are either crossbar or bus based architectures. In fact, while MINs were widely studied in the literature, we think that the use of MINs as an





FIG. 1 – A topological classification of MINs

FIG. 2 – A block diagram of $\Omega(16, 4)$

interconnection medium for SMP architectures can improve their performance. This might be the subject of future studies in the field.

The remainder of this paper is organized as follows : In section 2 we present a topological classification of MINs before a rapid reminder of the two example architectures to be used in the paper : the Omega network and the MCRB network. This is followed by the explanation of the evaluation methodology to be used. Section 5 presents the comparison and evaluation results before ending the paper with the conclusion.

2 Topological Classification of MINs

A MIN can be defined as a [SH94] network used to inter-connect a group of N inputs to a group of M outputs using several stages of small size SEs followed (or leded) by stages of wire links.

We propose in figure 1 a topological classification of MINs. We give in the following some important definitions related to this classification.

Definition 1 A banyan[GL73] MIN is a MIN having the property of the existence of one and only one path between each source and destination.

Banyan MINs might have the *delta property* or not. Delta networks, proposed by Patel[Pat81], are built of $a \times b$ crossbars. Let o_i ; $i = 0, 1, \dots, b - 1$ be an output of index i of a crossbar. If an input of a crossbar in stage j is connected to an output o_i of another crossbar in stage $j - 1$, then all its other inputs must be connected to outputs of the same index i of crossbars in the previous stage. We propose the following mathematical translation of the delta property.

Definition 2 For a Banyan MIN of size N and degree r^1 , suppose that the switch's inputs and outputs are presented to the base r , i.e. in the form d_0, d_1, \dots, d_{r-1} . Let the inputs and outputs of the SEs in the network have the same indexes then digits d_0 of all inputs of a switch must be equal. A network or a stage having this characteristic is called to be having the *Delta property*.

¹In this paper Network(N, r) will present a MIN of size N and degree r .

SEs in delta networks, which are banyan networks, are digit-controlled crossbars. Digit-controlled crossbars are controlled by digits of the message's control sequence. We call a control sequence the succession of digits representing the path to be taken by the message through the MIN. Usually, in Delta MINs this control sequence is a representation of the destination to the base r .

In fact, a network having the Delta property possesses some kind of regularity so that the network's routing algorithm can be simple and well defined [KS86].

Definition 3 *We call an over-sized MIN of size N a banyan Delta MIN composed of more than one copy of a Delta MIN gathered together by an interconnection stage having the Delta property.*

Non-banyan networks can be constructed either by the *augmentation* of a banyan network or by the construction of a *multipath* network such as the Clos network [Clo53].

Definition 4 *Augmenting [KS83] a MIN is the procedure of adding links and/or switches to a banyan network in order to make it a multipath network "without sacrificing much of its structure" [KS83].*

3 Case Studies

Our goal in this paper is to evaluate the effect of MINs degrees on their performance. The study will be applied on Delta and over-sized Delta MINs. Delta networks will be presented by its sub-class Omega network for comparison purposes later in this paper. The MCRB network will be the example used to study over-sized Delta networks.

3.1 Omega Network

Omega network, first defined by Lawrie [Law75], is a subset of the delta networks family proposed by Patel [Pat81], which is a bit-controlled interconnection networks family. In fact Lawrie's definition in [Law75] forms a special case of shuffle-exchange networks [HB89]. By example, omega network's size is not inevitably a power of 2. Furthermore, this kind of interconnection networks might be built using q-shuffle [Pat81] wiring stages so that crossbars of sizes different than 2 can be used to build Omega networks. Figure 2 shows an Omega(16,4) network.

3.2 Example of Over-sized Delta Networks : The MCRB Network

The MCRB topology, defined by Kechadi in [Kec02], is a dynamic multistage implementation based on the static chordal ring topology [PK99]. For complexity reasons, described in [ADKS02], we give here a definition a bit different of the one given by Kechadi.

Proposition 1 An $MCRB(r^n, r)$ network is a MIN built of $r \times r$ SEs and contains n stages of r^n (SEs) each. Let SE_{ij} be the switching element j of the stage i of the $MCRB(r^n, r)$, then SE_{ij} is connected to SE_{i-1, k_d} such that $k_d = (j + d r^i) \bmod N$, for $0 \leq i \leq N - 1, 1 \leq j \leq r - 1$, and $0 \leq d \leq r - 1$.

As an example, the configuration of the $MCRB(8, 2)$ is shown in Figure 3.




FIG. 3 – $MCRB(8, 2)$ Network

In [ADS03] it is proved that the MCRB network is a special case of Delta networks. We call this special case over-sized Delta networks.

4 Evaluation Methodology

In [ADS03] a methodology based on performance measures which defines a systematic decision making mechanism for choosing more suitable MINs for a multiprocessor system was presented. This methodology is used in this paper to evaluate the effect of the network's degree on its performance. We will limit our study to three performance evaluation factors knowing that the proposed methodology is general and that it is easy to add other factors chosen to evaluate the performance of a MIN.

Complexity is a quantitative term related to the *cost*. The evaluation of a system, whether it is hardware or software or both, needs a full study of the cost to be paid in order to build and implement it. Thus, the cost must be calculated in space and time terms. Complexity is not a stand alone parameter; when evaluating a system, many other performance factors such as throughput and different system components depend on it.

4.1 Integration complexity

While studying a MIN, the first evaluation to do is its hardware complexity. The hardware complexity of a MIN can be calculated by two means : the number of connection points

and the number of connections or wires needed to construct the MIN. Liu[Liu90] defines the hardware complexity of a MIN as the maximum of the two means. The hardware complexity of a MIN in term of crosspoints is equal to the total number of crosspoints of all crossbars used to build it. The complexity in terms of connections is the sum of links or wires in all stages.

Definition 5 *consider a MIN of size N and degree r , that has X stages of x SEs each. The stages are connected with Y inter-stages links. The integration complexity of the MIN will be defined as $C = \max(r^2 Xx, Yr)$.*

4.2 Temporal complexity

In spite of criticisms[Dun90], Flynn's taxonomy of functional environments for parallel architectures[Fly72] seems to stay the most accepted one. Basic performance criteria in SIMD environments are different of those in MIMD ones[All94]. When studying routing capacity of MINs, throughput is the important performance factor in MIMD environments, while in SIMD environments, the important criterion is the network's permutation capacity.

4.2.1 Throughput

This is defined as the number of messages delivered to their destinations per unit of time [Mer91, Pat81]. Many analytical studies of MIN's throughput can be found in the literature[Pat81, KS83, SH87]. Simulation is used frequently when more realistic results are needed. It allows more flexibility in network characterization in order to make it possible to analyze real-world and popular communication patterns. In fact, to study the throughput of an unbuffered network, messages leave sources to their destinations and in the case of a conflict, only one message goes through and the others are discarded. The throughput is calculated as the number of messages arrived to their destinations over a certain number of trials.

Definition 6 *In an unbuffered MIN, We define the throughput as the number of messages delivered to their destination per unit of time knowing that only one message goes through when more than one message assigned the same interconnection source. All other messages are discarded.*

4.2.2 Network's Latency

Another important performance parameter is the network's latency, which is defined below. The network latency analysis depends directly on the maximum number of cycles needed to route a certain number of permutations to their destinations. We use the same previously explained simulation to measure the analyzed MINs' latency.

Definition 7 *The Latency of a MIN is defined by the number of network cycles needed for all messages of a permutation to arrive to their destinations. This is referred to as the network's universality[SH94].*

4.3 Universal Performance Factor

Here we explain how many performance factors can easily be combined in order to get a universal performance evaluation factor. The above defined factors will serve as examples to apply our proposed evaluation methodology.

It is well known that the evaluation and the comparison of interconnection networks are not easy tasks. This is due to the very big number of criteria and factors that must be evaluated, the importance of every factor and its importance relatively to the other factors[LAS97].

We will suppose that the importance of the factors is a designing choice, i.e. we will suppose in our study that the performance factors to be evaluated are chosen. In general, performance evaluation factors belong to one of two major groups : factors to be maximized and factors to be minimized. We will call the group of factors to be maximized p^{max} and the one of factors to be minimized p^{min} , so that $p^{max} = \{p_1^{max}, p_2^{max}, \dots, p_k^{max}\}$ and $p^{min} = \{p_1^{min}, p_2^{min}, \dots, p_l^{min}\}$ where k is the number of factors to be maximized and l is the number of factors to be minimized. We define our universal performance factor as :

$$UPF = \sqrt{\sum_{i=1}^k (p_i^{min})^2 + \sum_{j=1}^l \frac{1}{(p_j^{max})^2}} \quad (1)$$

Using equation 1 and having two networks μ_1 and μ_2 with UPF_1 and UPF_2 in order as UPFs, we can say that if $UPF_1 < UPF_2$ then μ_1 is more powerful than μ_2 .

In order to clarify the idea behind this definition of a performance factor we will study an example of two MINs performance evaluation with only two performance factors and let them both be factors to be minimized. Let us plot on a 2 dimensions space the performance factors of the two MINs μ_1 and μ_2 . As an example, let figure 4 be a representation of this plot. Let p' and p'' be the two factors to be evaluated and let $p1'(p2')$ and $p1''(p2'')$ be the calculated values of these terms for $\mu_1(\mu_2)$.




FIG. 4 – An example of the use of the UPF factor

It is clear from figure 4 that the performance of μ_1 is better than this of μ_2 as for the two terms μ_1 gives smaller results than those of μ_2 . Now note that the UPF is the representation

of the distance between the 0 point and the point representing the performance of the network. Knowing this we notice that the less the UPF is, the better is the network.

This example can very easily be generalized on $k + l$ dimensions, which gives the formula of equation 1.

In fact when using calculated values to compare many performance factors a problem is to be studied; this is the difference of the *weight* of the different factors. In other words, when one or more values are very big comparing to the other studied factors, the comparison returns to be a comparison of this very big factor or these very big factors only. In order to solve this problem, values can be normalized to a certain value, which might be the maximum value, or better still, the average values for all the factors. We call \bar{p} the average of the values of the elements of the group p . This leads to modify equation 1 to the following equation.

$$UPF = \sqrt{\sum_{i=1}^k \left(\frac{p_i^{min}}{\bar{p}_i^{min}} \right)^2 + \sum_{j=1}^l \left(\frac{p_j^{max}}{\bar{p}_j^{max}} \right)^2} \quad (2)$$

Formula 2 can be further improved by including the importance aspect. This can be done by multiplying every term in this last equation by a constant related to the performance term that we call the *weight* (w) of the factor. This gives the following final equation to calculate the UPF.

$$UPF = \sqrt{\sum_{i=1}^k w_i \left(\frac{p_i^{min}}{\bar{p}_i^{min}} \right)^2 + \sum_{j=1}^l w_j \left(\frac{p_j^{max}}{\bar{p}_j^{max}} \right)^2} \quad (3)$$

From now on all performance factors are normalized values and all factors' weights will be one.

5 Performance Evaluation

In this section, results of the application of our methodology on several MINs are presented. Note that two stages MCRB networks will not be treated in this paper as they non-blocking networks. Moreover, this kind of MCRB networks is more complex than the crossbar. This point was studied in a previous publication[ADKS03].

5.1 One dimension evaluations

We mean by one dimension evaluations those considering only one performance factor.

5.1.1 Latency

Concerning latency, results that we got correspond to those found by Cheemalavagu and Malek[CM82]. Figure 5 shows that the latency of networks of degree 4 is better than those of degree 2 and 8 for 64 size SW-banyans. In fact, this is not true for another banyan special

case Delta network, which is the over-sized network. Figure 6 shows that, always, for the over-sized Delta network less latency is obtained with crossbars of bigger degrees. The same figure shows that there is an optimal value of the Omega network's degree so that it has the least latency. This can be explained as follows : It is evident that bigger size crossbars can pass a larger number of permutations of the same size. However, the permutation capacity of a crossbar might be limited by the great number of requests coming to its inputs. The probability of a conflict to occur in a 2 degree Delta network is greater than this in a bigger degree one. This can be demonstrated by simple calculation using, by example, Patel's formula [Pat81] to calculate the probability of message arrival in a Delta network. In fact, for a certain degree value, this conflict probability reaches an important level so that less messages can pass, which implies a larger latency. This explains the existence of the optimal value for Delta networks. On the other hand, as the number of crossbars in an over-sized Delta network is r times the number of crossbars in a Delta network, this limit is not reached and the universality of the network is always better when using crossbars of bigger sizes.




FIG. 5 – Latency of 64 size networks as a function of the degree

5.1.2 Throughput

Figure 7 shows the throughput of different sizes networks as a function of the degree. The figure shows that for over-sized Delta networks, throughput is bigger when larger size crossbars are used. On the other hand, Delta networks have, once again, an optimal degree value, which is 2 for size 64 networks, 2 and 4 for size 1024 and 4 for 4096 size network.

Note that when a very big number of messages arrive to the first stage of a network, using our approach to calculate the throughput, a considerable number of messages are discarded. No conflicts can occur on the last stage as the destinations are groups of permutations. This




FIG. 6 – Latency as a function of the degree




FIG. 7 – Throughput as a function of the degree

means that conflicts can occur only on a limited number of stages in large degree MINs which might explain the slight augmentation of the the throughput of $\Omega(4096,64)$ as related to $\Omega(4096,16)$.

5.2 Two dimensions evaluations

Here, two performance factors are evaluated simultaneously.

5.2.1 Complexity and throughput UPF

Comparing the UPFs of several MINs regarding only the complexity and the throughput means that the latency of the networks is not an important factor to be evaluated. In other words, the betterness of the network is judged by a small complexity and a high throughput only, which are to be evaluated at the same time. What is really important to note in figure 8 is the complexity effect on the performance of the over-sized Delta networks. This will be further cleared in the next section where we study the universality-throughput UPF.




FIG. 8 – Complexity-throughput UPF as a function of the degree

Note also the approach of the performance values of the Delta network to the optimal degree value. For small size networks, small degrees give the best results, while for networks of size 1024, crossbars of size 2 and 4 give the same best performance. However, with 4096 size network the optimal degree value is 4.

5.2.2 Universality and throughput UPF

Here, the cost of the MIN is not an important factor to be evaluated. So, we are looking for the MIN that gives the best latency AND throughput at the same time regardless to the complexity. In fact, MCRB networks show a very good performance when complexity is not considered. This is why their universality-throughput UPF is considerably smaller when using bigger sizes crossbars. Remember that smaller UPF means better performance. Note that complexity plays, when considered, an important role in degrading the network's performance as related to the complexity-throughput UPF case. This is because over-sized Delta networks' complexities augment very rapidly with their degrees.




FIG. 9 – Latency-throughput UPF as a function of the degree

On the other hand, we see once again the optimal value aspect when studying the Delta network. Note also the improvement of the UPF for the Omega(4096,64) relatively to the Omega(4096,16). This is due to the improvement of the throughput of this network, previously explained.

5.3 Inclusive UPF

Figure 10 presents the comparison and evaluation of a number of networks as related to the three factors we are considering in this paper as examples, i.e. complexity, universality and throughput.

Once more, we can observe the important effect of the complexity on the overall system performance. This can be easily seen by comparing figures 10 and 8, in which curves have almost identical shapes (the shapes are totally different when complexity is not considered, see figure 9). What is remarkable in figure 10 is that one of the over-sized Delta cases,




FIG. 10 – Inclusive UPF as a function of the degree

which is when the size of the network is 4096, has an optimal degree value which does not correspond to the biggest one. The best case for 4096 size over-sized Delta networks, when considering all three performance factors, is when the degree is equal to 4.

6 Conclusion

In this paper, we used a comparison and evaluation methodology applying simple heuristics to compare networks with different degrees. The main interest of the paper was to evaluate the network's degree effect on its performance. Our observations showed that an optimal degree value exists for Delta MINs. However, in general, for bigger degrees MCRB networks better performance is gotten. An exception is noted to this rule when network's complexity becomes really important and all performance factors are considered. In this case the over-sized Delta network can even approach the Delta performance network optimal value aspect.

This study of MINs is aimed to be a basis for future study of the use of MINs in SMP machines.

Références

- [ADKS02] A. Ch. Aljundi, J.-L. Dekeyser, M.-T. Kechadi, and I. D. Scherson. Comparative simulations and performance evaluation of mcrb networks using multidimensional

- queue management. In *Proc. Int'l Symp. on Performance Evaluation of Computer and Telecommunication Systems*, pages 288–296, San Diego, USA, July 2002.
- [ADKS03] A. Chadi Aljundi, Jean-Luc Dekeyser, M-Tahar Kechadi, and Isaac D. Scherson. A study of an evaluation methodology for unbuffered multistage interconnection networks. In *Proceedings of 17th International Parallel and Distributed Processing Symposium, IPDPS'03*, Nice, France, April 2003.
- [ADS03] A. Ch. Aljundi, J. L. Dekeyser, and I. D. Scherson. An interconnection networks comparative performance evaluation methodology : Delta and over-sized delta networks. To be presented in : PDCS'03, Reno, USA, Aug. 2003.
- [Al194] B. D. Alleyne. *Methodologies for Analysis and Design of Data Routers in Large SIMD Computers*. PhD thesis, Princeton Univ., June 1994.
- [Clo53] C. Clos. A study of non-blocking switching networks. *Bell system tech. journal*, 32(2) :406–424, Mar. 1953.
- [CM82] S. Cheemalavagu and M. Malek. Analysis and simulation of banyan interconnection networks with 2x2, 4x4 and 8x8 switching elements. In *Proc. Real-Time Syst. Symp.*, pages 83–89, 1982.
- [CSG99] D. E. Culler, J. P. Singh, and A. Gupta. *Parallel Computer Architecture (A Hardware/Software Approach)*, chapter Interconnection Network Design. Morgan Kaufmann Publishers, 1999.
- [Dun90] R. Duncan. A survey of parallel computer architectures. *IEEE Computer*, 23(2) :5–16, Feb. 1990.
- [Fly72] M.J. Flynn. Some computer organizations and their effectiveness. *IEEE Trans. Comput.*, C-21(9) :948–960, Sep. 1972.
- [GL73] G.R. Goke and G.J. Lipovski. Banyan networks for partitioning multiprocessor systems. In *Proc. 1st Annu. Symp. Comput. Arch.*, pages 21–28, 1973.
- [HB89] K. Hwang and F.A. Briggs. *Computer Architecture and Parallel Processing, 5th printing*. McGraw-Hill series in computer organization and architecture. McGraw-Hill International Editions, 1989.
- [KD97] M. T. Kechadi and J.-L. Dekeyser. Analysis and simulation of an out-of-order execution model in vector multiprocessor systems. *Parallel Computing*, 23 :1963–1986, 1997.
- [Kec02] M. T. Kechadi. Mcrb : A new interconnection network for multiprocessor systems. In *Misc. Papers, CD-ROM of the 2002 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'02)*, ISBN : 1-892512-39-4, Las Vegas, USA, June 2002.
- [KS83] C. P. Kruskal and M. Snir. The performance of multistage interconnection networks for mutliprocessors. *IEEE Trans. Comput.*, C-32(12) :1091–1098, Dec. 1983.
- [KS86] C.P. Kruskal and M. Snir. A unified theory of interconnection network structure. *Theoretical Computer Science*, 48 :75–94, 1986.

- [LAS97] K.J. Liszka, J.K. Antonio, and H.J. Siegle. Problems with comparing interconnection networks, is an alligator better than an armadillo? *IEEE Concurrency*, 5(4) :18–28, October-December 1997.
- [Law75] D. A. Lawrie. Access and alignment of data in an array processor. *IEEE Trans. Comput.*, C-24(12) :1145–1155, Dec. 1975.
- [Liu90] Y.-S. Liu. *Architecture and performance of processor-memory interconnection networks for MIMD shared memory parallel processing systems*. PhD thesis, New York University, 1990.
- [Mer91] A. Merchant. *Analytical Models for the Performance Analysis of Banyan Networks*. PhD thesis, Stanford University, 1991.
- [Pat81] J. H. Patel. Performance of processor-memory interconnections for multiprocessors. *IEEE Trans. Comput.*, C-30(10) :771–780, Oct. 1981.
- [PK99] B. Parhami and D.-M. Kwai. Periodically regular chordal rings. *IEEE Transactions on parallel and Distributed Systems*, 10(6) :658–767, Jun. 1999.
- [SH87] T.H. Szymanski and V.C. Hamacher. On the permutation capability of multistage interconnection networks. *IEEE Trans. Comp.*, C-36(7) :810–822, Jul. 1987.
- [SH94] T.H. Szymanski and V.C. Hamacher. On the universality of multistage interconnection networks. In *Interconnection Networks for High-Performance Parallel Computers*, pages 73–101. IEEE Computer Society Press, 1994.
- [Sie77] H.J. Siegel. Analysis techniques for simd machine interconnection networks and the effects of processor address masks. *IEEE Trans. Comput.*, C-26(2) :153–161, Feb. 1977.
- [SY94] I. D. Scherson and A. S. Youssef. *Interconnection networks for high-performance parallel computers*. IEEE computer society press, 1994.
- [WB72] W.A. Wulf and C.G. Bell. C.mmp- a multi-mini-processor. In *AFIPS Proc. Fall Joint Computer Conf.*, pages 765–777, 1972. NL.



Unité de recherche INRIA Futurs

Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399