An Interconnection Networks Comparative Performance Evaluation Methodology: The Case of Delta and Over-Sized Delta Multistage Interconnection Networks

Ahmad Chadi Al Jundi, Jean-Luc Dekeyser
Laboratoire d'Informatique Fondamentale de Lille
Universite des Sciences et Technologies de Lille
59650 Villeneuve d'Ascq
France
{aljundi,dekeyser}@lifl.fr

Isaac D. Scherson
Dept. of Information and Computer Science
University of California, Irvine
Irvine, CA 92697
USA
isaac@uci.edu

Abstract

Interconnection network performance is a key factor when constructing parallel computers. The choice of an interconnection network used in a parallel computer depends on a large number of performance factors which are very often application dependent. We propose a performance evaluation and a comparison methodology. This methodology is applied on a new class of interconnection networks (Over-Sized Delta Network) and on the Omega network, and will be used in future work in order to evaluate the use of multistage interconnection networks as an intercommunication medium in today's Symmetric Multiprocessors.
Keywords: Interconnection Networks, Parallel Architectures, Delta Networks, Banyan Networks, Performance Evaluation.

1 Introduction

Inter-linking processors and linking them efficiently to the memory in a parallel computer are not easy jobs because of conflicts that might happen when more than one communication task take place at the same time[10]. It is evident that a simple bus is not an enchanting solution as simple busses can not transfer more than one message simultaneously. On the other hand, a crossbar, which provides a full connection between all the nodes of the system is very complex, expensive, and hard to be controlled. Therefor, Interconnection Networks[22] were a good means used as intercommunication media in these parallel systems.

Multistage Interconnection Networks (MINs) are usually used in multi-processor and multi-computer parallel machines as an intercommunication medium between Processing Elements (PEs) 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. Kruskal and Snir[12] studied *augmentation* techniques that might be applied on banyan networks in order to improve their performance without much loss in simplicity. One of these techniques is the *d-replication* of the network, which is the use of d copies of a network of size N to connect N sources and N destinations.

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 are intending to investigate in a future work the use of MINs as an interconnection medium for SMP architectures.

In this paper we investigate the use of several copies of a Delta network of size N in order to interconnect N sources and N destinations without, practically, loosing the banyan characteristic of the network. We will call this technique over-sizing. In fact over-sizing a delta network results in a network defined in [11] called the MCRB network. The transfer procedure between the two networks is mathematically proved and then the performance of the resulted network is evaluated. We propose also a simple evaluating performance factor that might be used for all MINs and all *important*

performance factors. The importance and the relative importance of the factors are taken in consideration by the proposed factor. This study will serve in the future as a basis for the evaluation of the use of MINs in SMP machines.

The remainder of this paper is organized as follows: after the introduction we propose a topological classification of MINs. In section 3 some examples of MINs are studied in some detail and then we propose a MINs evaluation methodology in section 4. Some results of the application of this methodology on some example MINs are presented in section 5 before concluding the paper.

2 Topological Classification of MINs

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[8] 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[20], are built of $a \times b$ crossbars. Let o_i; $i = 0, 1, \ldots, b-1$ 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 $d_0, d_1, \ldots, d_{r-1}$. 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.*

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.

[1]In this paper Network(N,r) will present a MIN of size N and degree 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[13].

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[5].

Definition 4 *Augmenting[12] 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"[12].*

Kruskal and Snir studied in [12] two augmentation strategies: replication and dilation. The following two definitions are almost copied from Kruskal and Snir's paper.

Definition 5 *The d-dilation of a network G is defined to be the network obtained from G by replacing each edge (link) by d distinct edges.*

Definition 6 *The d-replication of a network G is defined to be the network consisting of d identical distinct copies of G. This definition will be explained later.*

3 Case Studies

Our goal in this paper is to evaluate the performance of over-sized Delta MINs. Their performance is to be compared to the performance of Delta MINs in order to realize the gain obtained by the over-sizing. As Omega network forms a sub-class of Delta networks, it will be presented and used later for the comparison as a Delta network. On the other hand, the equivalence between MCRB and over-sized networks is to be proved and the MCRB network will be compared with Delta MINs. At the end of this section we will reminde of the definition of replicated MINs and show the difference between them and over-sized networks.

3.1 Omega Network

Omega network, first defined by Lawrie[14], is a subset of the delta networks family proposed by Patel[20], which is a bit-controlled interconnection

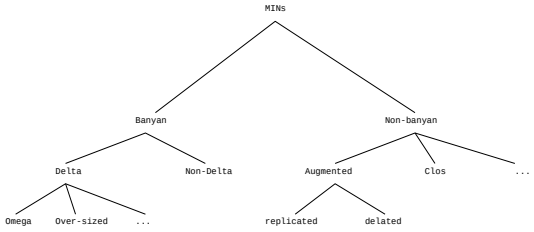

Figure 1: A topological classification of MINs

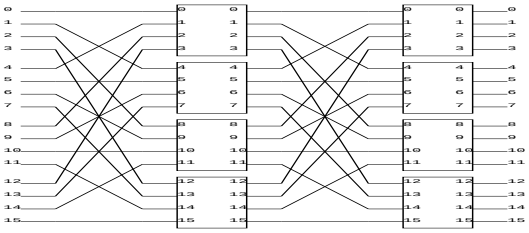

Figure 2: A block diagram of $\Omega(16,4)$

networks family. In fact Lawries definition in [14] forms a special case of shuffle-exchange networks[9]. 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[20] wiring stages so that Omega network might be built using crossbars of degrees different of 2. Figure 2 shows an Omega(16,4) network..

Using this definition, Omega network might be built using crossbars of degrees different of 2. 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 [11], is a dynamic multistage implementation based on the static chordal ring topology[19]. For complexity reasons, described in [2], we give here a definiton a bit different of the one given by Kechadi.

Proposition 1 *An $MCRB(r^n, r)$ nework is a MIN built of $r \times r$ SEs and containes 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.

MCRB networks can be seen as a special case of Delta networks built by replacing every switch in the Delta network with r switches of the same size. We will explain the procedure of getting the corresponding Delta network of an MCRB network in the following.

Proposition 2 *Let $\mu(N,r)$ be an MCRB network of size N and degree r. In order to derive the corresponding Delta network, switches of distance r in μ must be regrouped in order to form a Delta(N,r) network. Moreover, the r corresponding Delta networks can be connected by a last stage. The overall resulting network is a Delta network.*

Proof. The proof will be discussed in two steps: we prove in step one that the topology of the first $n-1$ stages of an MCRB network is equivalent to the topology of r Delta networks. In the second step we prove that relying the r delta networks with a stage having the topology of stage 0 of an MCRB network results in a Delta network of size $N \times r$.

Step 1. From proposition 1 we see that a switch j in a stage i is connected to switches of which indexes are $j + dr^i$ in stage $i+1$. In other words we can say that switch x in stage i is connected to switches of which the indexes are of a difference equal to multiplications of $r^i; i > 0$ which yields in indexes having the same digits to the base r and so corresponding to the definition 2.

Step 2. As the MCRB network is divided into r identical Delta networks, it is clear that if a switch in the last stage receives a link from an output of index i, all the other inputs will receive links from outputs of index i.

Applying this procedure on an MCRB(8,2) shown on figure 3 gives the equivalent Delta network of figure 4.

Definition 7 *In [21] two MINs are defined to be equivalent if and only if they can realise the same permutations by adding a wired permutation stage to one of them.*

Proposition 3 *Every over-sized Delta MIN has an MCRB equivalent network.*

Proof. As every MCRB network is composed of r Delta networks and a Delta linking stage, we can say: For all values of N and r for which r is a power of 2 and $\log_r N = n\ \exists\ \Delta(N/r, r) \& \Delta_s$ where Δ_s is a linking stage having the Delta property, which can be combined as an MCRB network.

3.3 Comparing over-sized MINs to Replicated MINs

Definition 8 *The d-replication of a MIN of size N is defined by the use of d identical copies of the original*

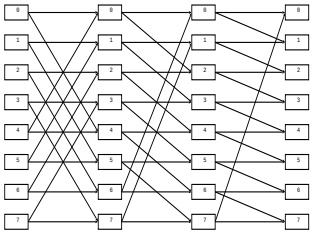

Figure 3: $MCRB(8,2)$ Network

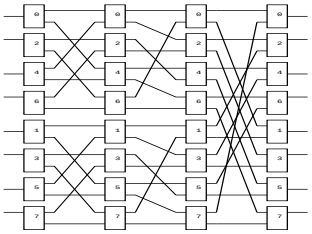

Figure 4: The Delta network equivalent to MCRB(8,2)

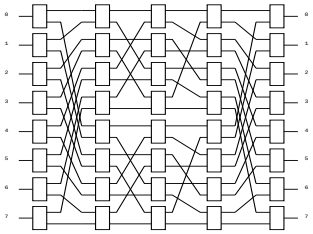

Figure 5: A 2-replicated Delta(8,2) MIN

MIN in order to connect N sources and N destination. Figure 5 shows a 2-replicated Delta(8,2) MIN.

Comparing the two networks of figures 4 and 5 we notice that over-sized MINs are a mid-step between banyan delta networks and replicated "multipath" networks. In fact, replicating a delta network consists of linking the copies of the network with two shuffle linking stages: one for the inputs and another for the outputs. It is easy to prove that shuffle linking possesses the *delta property*. This means that the only difference between replicated and over-sized MINs is that in the second one there is only one additional stage in order to link the several copies of the Delta network. On the other hand, while in replicated MINs we gain the multipath property relatively to banyan Delta networks, what we gain in over-sized networks is $1/r$ less work load distributed over all switches. Note also that over-sized networks are less complex than replicated networks.

We are interested in this paper in the performance evaluation of this kind of over-sized networks. From now on, we will refer to over-sized Delta networks As MCRB networks.

4 Evaluation Methodology

One of the main goals of this paper is to provide a methodology which defines a systematic decision making mechanism for choosing more suitable MINs for a multiprocessor system. This methodology is based on performance measures. 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[17] 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 9 *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* $C = max(r^2Xx, Yr)$.

4.2 Temporal complexity

In spite of criticisms[6], Flynn's taxonomy of functional environments for parallel architectures[7] seems to stay the most accepted one. Basic performance criteria in SIMD environments are different of those in MIMD ones[4]. 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 [18, 20]. Many

analytical studies of MIN's throughput can be found in the literature[20, 12, 23]. Simulation is used frequently when more realistic results are needed. It allows more flexibility in network characterization in order to make it possible network 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 10 *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 Permutation capacity

Permutation capacity [1] refers to the MIN's capacity to route permutation messages. We mean by permutation messages groups of messages of which destinations are permutations of all inputs. In certain cases, requests destinations are permutations of all or a subset of a system's memory modules [23].

In order to study the permutation capacity of a MIN, one might try to route a certain number of random permutations on the network and calculate the number of cycles needed to route all permutation messages to their destinations. Analytical studies [15] as well as our experience proved that routing random permutations is not efficient for the permutation capacity study. Therefore, frequently used permutations [15] must be used for such an analysis. To establish this comparative study, a certain number of these permutations can be routed under the assumption that the accumulated number of permutations routed per cycle can be a comparison factor between them.

Definition 11 *Permutation capacity of a MIN is a factor that measures the number of permutation messages that can arrive totally to their destinations in a certain number of interconnection cycles.*

To calculate the permutation capacity of a MIN we simulate routing permutations of one family of frequently used permutations, which is the BPC (Bit Permute/Complement) family.

Definition 12 *[4] A BPC permutation of a set of integers is another set of integers whose binary representations are permutations of the binary representations of the first set. Some bits of each representation of the result might then be 1's complemented.*

4.2.3 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 13 *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[24].*

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[16].

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 $p^{max} = \{p_1^{max}, p_2^{max}, \ldots, p_k^{max}$ and $p^{min} = \{p_1^{min}, p_2^{min}, \ldots, p_l^{min}\}$ 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:

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

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 6 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 $\mu_1(\mu_2)$.

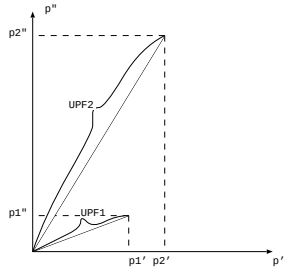

Figure 6: An example of the use of the UPF factor

It is clear from figure 6 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 fcator 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 $\overline{p}$ the average of the values of the elements of the group p. This leads to modify equation 1 to the following equation.

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

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.

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

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

5 Over-sized Delta networks' performance

Two and three dimensions evaluations are to be presented in this section. Networks of different sizes and degrees are to be tested. Note that *one dimension* evaluations are presented in previous publications in [2] and [3]. Note that not every composition of factors is available to be used in order to compare MINs. By example, the use of the complexity with the latency does not represent an acceptable test as MINs of small size are not complex and do not need a lot of cycles to rout the simple permutations that they can do.

5.1 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. On the other hand, the betterness of the network is judged by a small complexity and a high throughput (which is equivalent to small values on the y axis of figure 7).

Figure 7 shows a comparison of different MCRB and Omega MINs. Note that it is clear that MCRB(1024,4), MCRB(512,8), and MCRB(1024,2) are too complex relatively to the throughput that they provide. On the other hand, the throughput of Omega(512,8) network is very small relatively to its cost. All other compared networks of figure 7 are very close regarding their UPFs. The best, i.e. the one that has the smallest UPF is the MCRB(256,4).

5.2 Latency 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 of the complexity. In fact MCRB networks show a very good performance when

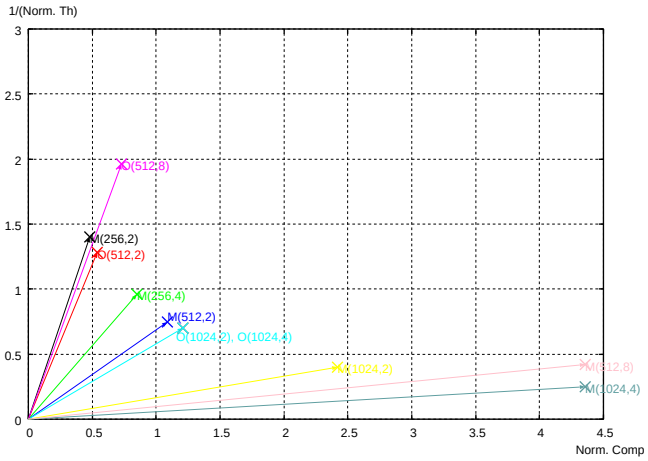

Figure 7: Comparing some MCRB and Omega MINs regarding their throughputs and complexities

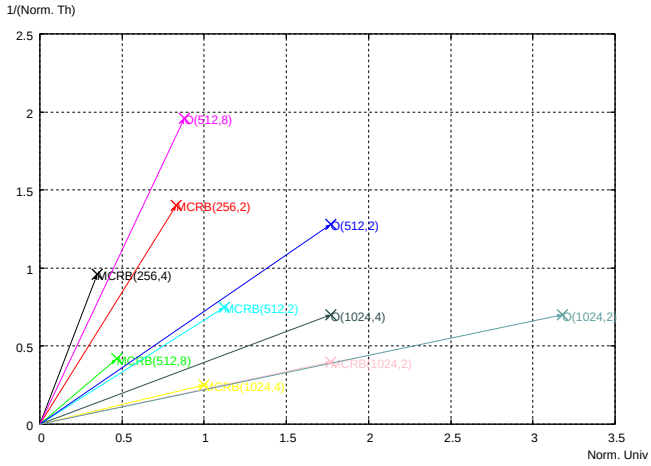

Figure 8: Comparing some MCRB and Omega MINs regarding their throughputs and Universalities

complexity is not considered. MCRB(512,8) (which has the biggest complexity among all studied MINs, is the most powerful network. MCRB(256,4) and MCRB(512,2) which were among the best networks when considering complexity and throughput are still among the best MINs in the case of considering latency and throughput.

5.3 Inclusive UPF

The most general case regarding the factors to be studied in this paper is the case in which all the studied factors, i.e. complexity, throughput, and latency, are taken in consideration.

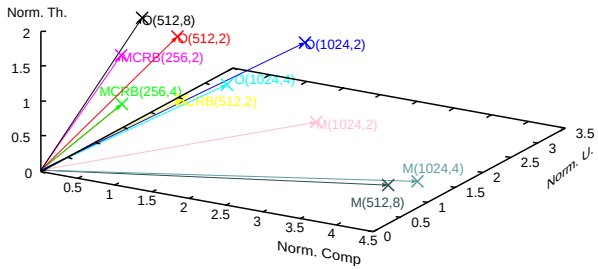

Figure 9: Comparing some MCRB and Omega MINs regarding their throughputs, Universalities and Complexities

Note that although MCRB(512,8) gives the best results when considering latency and complexity, the cost that is to be paid to get this performance is relatively high, and so, it is not considered as a very good choice among the others, i.e. less complex but not as good MIN.

On the other hand, MCRB(256,2) showed a relatively good performance in all studied cases and thus it can be considered as the best MIN among the tested ones.

6 Conclusion

In this paper we studied the over-sizing of Delta MINs, which is a possible performance improving strategy of Delta MINs. We proposed as well an evaluation and comparison methodology of MINs. This methodology was applied on over-sized Delta and Delta MINs. Over-sized Delta MINs seem to be more powerful but this is due to a higher complexity relatively to Delta networks. The proposed methodology is to be used in future work in order to exterminate the use of MINs in today's SMP machines.

References

[1] D. P. Agrawal. Graph theoretical analysis and design of multistage interconnection networks. *IEEE. Trans. Comp.*, C-32(7):637–648, Jul. 1983.

[2] 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.

[3] 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.

[4] B. D. Alleyne. *Methodologies for Analysis and Design of Data Routers in Large SIMD Computers.* PhD thesis, Princton Univ., June 1994.

[5] C. Clos. A study of non-blocking switching networks. *Bell system tech. journal*, 32(2):406–424, Mar. 1953.

[6] R. Duncan. A survey of parallel computer architectures. *IEEE Computer*, 23(2):5–16, Feb. 1990.

[7] M.J. Flynn. Some computer orgabizations and their effectivenes. *IEEE Trans. Comput.*, C-21(9):948–960, Sep. 1972.

[8] G.R. Goke and G.J. Lipovski. Banyan networks for partitioning multiprocessor systems. In *Proc. 1st Annu. Symp. Comput. Arch.*, pages 21–28, 1973.

[9] 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.

[10] M. T. Kechadi. *Un Modle de Fonctinnement Dsordonn Pour les Systmes Multiprocesseurs Pipelines Vectoriels Mmoire partages (Dfinition, Modlisation et Proposition d'Architecture).* PhD thesis, Universit des Sciences et Technologie de Lille, Laboratoire d'Informatique Fondamental de Lille, Mar. 1993.

[11] 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.

[12] C. P. Kruskal and M. Snir. The performance of multistage interconnection networks for mutliprocessors. *IEEE Trans. Comput.*, C-32(12):1091–1098, Dec. 1983.

[13] C.P. Kruskal and M. Snir. A unified theory of interconnection network structure. *Theoretical Computer Science*, 48:75–94, 1986.

[14] D. A. Lawrie. Access and alignment of data in an array processor. *IEEE Trans. Comput.*, C-24(12):1145–1155, Dec. 1975.

[15] J Lenfant. Parallel permutations of data: A benes network control algorithm for frequently used permutations. *IEEE Trans. Comp.*, C-27:637–647, July 1978.

[16] 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.

[17] Y.-S. Liu. *Architecture and performance of processor-memory interconnection networks for MIMD shared memory parallel processing systems.* PhD thesis, New York University, 1990.

[18] A. Merchant. *Analytical Models for the Performance Analysis of Banyan Networks.* PhD thesis, Stanford University, 1991.

[19] B. Parhami and D.-M. Kwai. Periodically regular chordal rings. *IEEE Transactions on parallel and Distributed Systems*, 10(6):658–767, Jun. 1999.

[20] J. H. Patel. Performance of processor-memory interconnections for multiprocessors. *IEEE. Trans. Comput.*, C-30(10):771–780, Oct. 1981.

[21] I. D. Scherson. Orthogonal graphs for the construction of a class of interconnection networks. *IEEE Trans. Parallel and Distributed Systems*, 2(1):3–19, Jan. 1991.

[22] I. D. Scherson and A. S. Youssef. *Interconnection networks for high-performance parallel computers.* IEEE computer society press, 1994.

[23] 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.

[24] 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.