

# Studying the effect of the design parameters on the interconnection network performance in NOWs\*

R. Martínez, J.L. Sánchez, F.J. Alfaro

Departamento de Informática

Universidad de Castilla-La Mancha, 02071 - Albacete (España)

{raulmm,jsanchez,falfaro}@info-ab.uclm.es

V. Chirivella

Departamento de Estadística e Investigación Operativa

Universidad Politécnica de Valencia, 46071 - Valencia (España)

vchirive@eio.upv.es

J. Flich

Departamento de Informática de Sistemas y Computadores

Universidad Politécnica de Valencia, 46071 - Valencia (España)

jflich@gap.upv.es

## Abstract

*With the increasing use of Network of Workstations (NOWs) as an alternative to huge parallel computers it has become essential to design high-performance interconnection networks for the communication between the nodes of these clusters.*

*A large number of studies have been carried out to achieve this objective. Most of them propose a new technique that affects one of the parameters that characterize the interconnection network. These techniques are completely new or inspired in the techniques previously used in multiprocessor systems. The impact of the proposal is studied (in most cases using simulation), and an analysis is made of the effect of the new technique over the system performance versus those currently in existence. In this kind of study most of the network parameters are fixed and usually only a few parameters are varied.*

*This paper presents a more general study of the interconnection network performance. This study consists in showing the effect of different design parameters over the network performance, and the interaction between them. This study would not be viable with the traditional techniques due to the number of simulations required. The alternative of the Experimental Design is used to carry out the study.*

---

\* This work was partly supported by the Spanish CICYT under Grant TIC2003-08154-C06 and JCC de Castilla-La Mancha under Grant PBC-02-008.

## 1. Introduction

During the last years, NOWs have become a cost-effective alternative to massively parallel computers. This is due, mainly, to the fact that high-end personal computers have dramatically increased their performance while exhibiting a relative low cost. At the same time, high-speed interconnection networks are available to efficiently connect all the processors in a NOW. The interconnection network plays a key role in the overall performance of the system. In this sense, the interconnection network technology has taken important steps forward, and is now able to provide high bandwidth and low latency [1, 8, 5].

This situation allows for many applications, which due to their requirements were limited to very specific systems, to be executed over new platforms like clusters. For instance, some of these applications are scientific computing, web servers, videoconferencing, multimedia servers, collaborative environments, distributed data bases, etc [4].

The fast advance in high-performance networks for clusters is partially due to the reuse of research that has been carried out for years in multiprocessor networks. Many techniques and ideas used in multiprocessor environments have been transferred to NOWs [3]. Nevertheless, this transfer is not always immediate or even possible. In some cases, the transfer can be made but at the expense of sacrificing some particular characteristics of the clusters (scalability, flexibility in the interconnection, expansibility, etc.).

The diversity of applications running over clusters force

networks to provide very high performance. There are applications, most of them parallel ones, which require communications with low latency, and there are many others, of a distributed nature, that also require high bandwidth. In this sense, in the literature, a large number of studies have been centered on increasing the performance of the interconnection network for clusters [2]. Some of these works focus on different parts/components of the network (network interfaces, switch architectures, links). Others, focus on adopting (or modifying the network in order to adopt) existing methodologies to the new environment. Finally, in many cases new strategies (routing algorithms, virtual channels utilization, etc) are proposed to obtain better results.

To sum up, in the last years, a great deal of effort has been put into research to consolidate the NOW environment as a suitable platform for running high performance distributed and parallel applications.

Many of the works related to interconnection networks follow a similar approach in the evaluation process. Typically, these works propose a new idea, concept or technique that allows increasing network and overall system performance. These ideas are then evaluated in a particular context defined by a set of network and system parameters. The results are typically contrasted against performance results without the use of the proposed technique and then, the benefits of the new idea are highlighted.

In order to get results, simulation of the entire network is commonly used. Unfortunately, a good network model requires a large number of parameters (i.e., switch architecture, flow control, packet sizes, traffic loads, routing algorithm, etc.) to be defined. Usually, the evaluation process consists basically in fixing a large set of parameters and only changing few of them (intentionally those that seem will influence more the performance when the proposed technique is used). Additionally, the number of different cases analyzed is not large, and thus, the number of values analyzed for each parameter is low. This way of approaching the performance evaluation may lead to results that are too tied to the fixed conditions, and thus, do not allow to estimate the possible interactions between parameters.

Taking this into account, it seems more reasonable to evaluate a new idea in a much more broad analysis that takes into account more parameters and interactions of different parameters. This broader analysis would reflect a more realistic view of the performance achieved by the new idea or concept. At first sight, in such a study, parameters are changed by turn while fixing the rest. Therefore, all the possible combinations of the different values of the parameters must be tested in order to achieve the required information, namely the effect of each parameter over the system performance, and the interaction between them. This procedure wastes time and money since it is not efficient neither effective to get the objective of evaluating the interconnection

network performance.

However, a much clever approach is to design an *experimental design*. The goal of experimental design is to obtain the maximum information with the minimum number of experiments. By doing this, the new idea or concept can be fully analyzed and detailed interactions with all the network parameters can be obtained with a reasonable amount of analysis. In this paper we will apply an experimental design to a particular network in order to fully analyze the interactions of the most critical parameters of the network.

The rest of this paper is organized as follows. Section 2 presents the experimental design technique. The approach to the performance evaluation is stated in Section 3, and the obtained results are shown in Section 4. Finally, some conclusions and future work are given in Section 5.

## 2. Design and Analysis of Experiments

In general, an *experimental design* will provide a set of experiments (number of experiments and value of parameters) for performance evaluation. As mentioned earlier, the goal of a proper experimental design is to obtain the maximum information with the minimum number of experiments. A proper experimental design also makes possible to separate out the effects of various design parameters that might affect the performance of the system being studied.

Once the experiment has been carried out, the analysis of the results determines if a design parameter has a significant effect, or if the observed difference is simply due to random variations caused by measurement errors and parameters that were not controlled.

### 2.1. Design of experiments

There are numerous varieties of experimental designs. The most frequently used are full factorial designs and fractional designs [6]. Before explaining each of them, several terms should be introduced. The *response variable* is the outcome of the experiment, generally the variable that measures the performance. Each design parameter that affects the response variable is called a *factor*. The values that a factor can assume are called its *levels*.

**2.1.1. Full factorial designs.** A full factorial design explores every possible combination of all levels of factors. A performance study with  $k$  factors, with the  $i$ th factor having  $n_i$  levels, requires  $n$  experiments, where

$$n = \prod_{i=1}^k n_i \quad (1)$$

The advantage of a full factorial design is that the effect of each factor and its interactions can be found. However, the main problem is the cost of the study. It may take too

much time (or money) to conduct all these experiments, especially when we take into account the possibility that each of them could be repeated several times. There are three ways to reduce the number of experiments: by reducing the number of levels for each factor; by reducing the number of factors; and by using fractional factorial designs.

The first alternative presents as its main disadvantage the fact that it could not be able to study all the desired levels of a quantitative factor, or the nature (lineal, quadratic) of the relationship between the response variable and a given factor. A full factorial design in which each of the  $k$  factors is used at two levels requires  $2^k$  experiments. This is a very popular design and is called  **$2^k$  design**. This alternative may be recommendable in an initial study where the goal is only to obtain the relative importance of each factor and not the nature of the relationship. After reducing substantially levels per factor, other factors can be tried.

The second alternative would need to restate the initial objectives of the problem resolution, and it may not suppose a substantial loss of information if the factors that are going to be left out have little effect on the response variable.

**2.1.2. Fractional factorial designs.** A fractional factorial design consists in a factorial design where an important number of experiments have been removed, at the cost of not being able to study some interactions of high order. For example, a full factorial design with 6 factors and 2 levels per factor would require 64 experiments (Equation 1), and it would allow us to estimate 6 simple effects<sup>1</sup>, 15 double interactions<sup>2</sup>, 20 triple, 15 quadruple, 6 quintuple and 1 sextuple, for a total of 63 effects. Many of these estimated effects, usually those of order higher than or equal to three, will be probably non-existent; this is because fractional factorial designs are appropriate in many cases.

As stated, a fractional factorial design consists in choosing only certain combinations out of all the ones possible (a full factorial design). For example, a  $2^{k-1}$  design would be a  $2^k$  design where a half of the experiments have been removed. In general, these are called  $2^{k-p}$  designs, allowing us to study  $k$  factors with  $2^{k-p}$  experiments.

To simplify the task of designing high fractional designs without mastering the theoretical basics, Genini Taguchi proposed the generalized use of some basic designs that he grouped in a set of tables and graphics [9]. These designs are called *orthogonal arrays*.

What makes it possible to calculate various effects and their contribution to the total variation is the concept of orthogonality. Given two factors,  $FI$  with  $I$  levels and  $FJ$

with  $J$  levels, it can be stated that the *simple effects of both factors are orthogonal if in the tests of the design each of the variants  $i$  of  $FI$  appears in the same proportion as the  $J$  variants of  $FJ$* . In a full factorial design the orthogonality is inherent, but a fractional factorial design must be carefully designed to maintain it.

## 2.2. Analysis of results

After the experiments have been designed and carried out, the useful information must be extracted from the results. To discover what effects are significant or not over the response variable the *Analysis of the Variance technique (ANOVA)* is used. When the importance of a specific factor has been determined, it is normal to use (according to the nature of the factor) the Least Significant Difference (LSD) intervals and the orthogonal contrasts to analyze in which way those factors affect the response variable.

**2.2.1. ANOVA** The study of the significance of the different factors and their possible interactions is carried out by elaborating the ANOVA table. The basic idea consists in splitting up the total variability observed in the results into a number of terms. Each term is associated to a different considered effect (simple effects, double interactions, etc.), plus one additional term associated with those effects that have not been considered and the possible error due to randomness, which is called the residual term. This variability is represented in the ANOVA table as the sum of the squares of differences between means. In the ANOVA table and for each factor, the fraction between the variation explained by each effect and the residual term is shown. If this fraction is high enough (taking into account an F of Snedecor distribution) the associated term is considered significant.

**2.2.2. LSD intervals and orthogonal contrasts** After determining the significance of an effect, the next step is to determine the nature of the relationship with the response variable. In the case of a qualitative factor, means can be compared, but by taking into account the LSD interval. The LSD interval is the approximate interval where the real value of the mean can be found with a certain level of confidence. If the factor having a significant effect is quantitative, orthogonal contrasts must be used. This technique consists in studying separately the effect of the different components (lineal, quadratic, etc.). These effects are analyzed in a similar way as in ANOVA. If a factor has  $i$  levels,  $i - 1$  components can be studied, but in general effects higher than quadratic are not significant.

1 A simple effect is the mean variability produced in the response variable when a single factor varies.

2 Interaction between two factors exists (double interaction) if the effect of one of them changes when the other varies its level. This concept can be generalized for superior orders.

### 3. A systematic approach to performance evaluation

To carry out this performance study, the methodology described in [6] has been used. This methodology suggests a certain number of steps that are fully developed in [7], but due to lack of space shown summarized form here.

#### 3.1. State goals and define the system

The goal of the study is to evaluate the performance of the interconnection network in a NOW according to different possible configurations. To simplify the study, the network consists of one switch attached to a certain number of hosts. The hosts will generate messages at a certain rate and messages will travel through the links to the switch, which will in turn send them to their destinations.

**3.1.1. Switch model** Figure 1 shows the switch model that will be used. There are input and output channels, an internal crossbar that provides connectivity between them, and a routing and arbitration unit that decides which output channel a message should be sent through when it arrives at the switch. Physical channels may be split into several virtual channels.

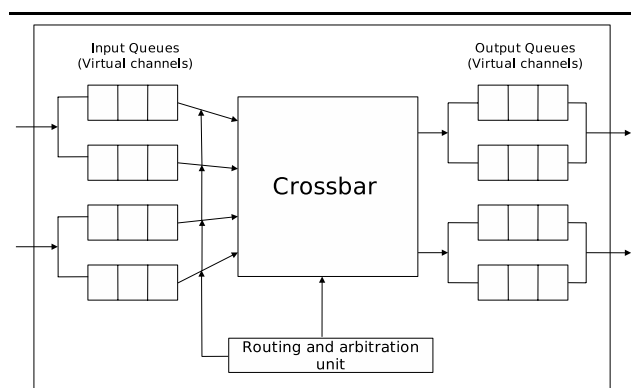


Figure 1. Switch model

#### 3.2. List services and outcomes. Select metrics

The goal of a network interface in a NOW is to efficiently transfer messages among the attached nodes. In this work, the chosen metric has been the throughput. The throughput is the maximum traffic accepted by the network, and is usually expressed as the number of received packets per node and time unit.

#### 3.3. List parameters. Select factors to study

The factors that have been chosen among all the possible parameters to carry out the study are shown in the following list:

- The factors that refer to the network itself are:
  - Flying time (TFLY): this is the time that the packet header takes to go from a host to a switch or vice versa.
  - Number of hosts (NHOST): this is the number of hosts connected to the switch.
- The factors that refer to the switch are:
  - Queue assignment policy (POLIC): this is the policy that decides which virtual channel corresponds with each message arriving to the switch.
  - Number of virtual channels (NVC): this is the number of storage queues associated with each port of the switch.
  - Crossbar organization (MUX): this indicates if the virtual channels corresponding to an input or output port are multiplexed or demultiplexed both at the input side or output side of the crossbar.
  - Routing table access (TABLE): this indicates the number of accesses that can be made to the routing table simultaneously. It ranges from one to infinite (parallel access) corresponding to the ideal case.
  - Routing time (TROUT): this is the time taken by the switch to obtain the output for a packet.
  - Arbitration time (TARB): this is the time taken by the arbitration algorithm to select an input and an output crossbar connection in order to forward a packet.
  - Bytes per cross (FCROSS): this is the number of bytes that traverse the crossbar for each connection in one cycle.
  - Size of queues (QSIZE): this is the size of the queues for each virtual channel.
- The workload factors are:
  - Message size (MESSIZE): this is the number of flits of the message.
  - Destination distribution (DDEST): this indicates the policy used to generate the message destinations.

Table 1 shows the different levels for the factors. It also specifies if the factors are considered to be design or noise factors, and if they are quantitative or qualitative. The design factors are those that must be fixed when the network

Factor	Levels	D/N	QN/QL
FCROSS	1, 2, 4	D	CN
DDEST	U, 1HS5%, 1HS10%, 1HS15%, 2HS5%, 2HS10%, 2HS15%	N	CL
LONGMEN	32, 164, 1024	N	CN
MUX	In&Out, In&-, -&Out, -&-	D	CL
NCV	1, 2/3/4, number of hosts	D	CN
NHOST	8, 16, 64	D	CN
POLIT	Aleatory, Most Empty Queue, VOQ	D	CL
TABLA	1, 2, 4, Parallel	D	CL
TAMCOLA	1, 4 messages	D	CN
TARB	1, 25, 50	D	CN
TFLY	8, 32, 64	D	CN
TROUT	1, 25, 50	D	CN

**Table 1. Factors considered. D/N stands for design/noise and QN/QL for quantitative/qualitative**

is physically made. The noise factors are those that affect the network performance but cannot be controlled. The levels for DDEST are the uniform distribution, and the hot-spot distribution with one or two hot spots and 5%/10%/15% of messages addressed to hot spots. The levels for MUX are the four possible combinations for multiplexing the virtual channels. The value for the medium level of NVC depends on the number of hosts in the network (NHOST), two virtual channels when there are eight hosts, three when there are sixteen and four when there are sixty-four.

### 3.4. Select evaluation technique. Select workload

Simulation has been chosen for performance evaluation because it will allow us to model a complex system with a large number of parameters to vary at a reasonable price.

Two message destination distributions have been used: uniform and hot-spots.

### 3.5. Design experiments

Taking into account the number of factors and levels considered in the study, and according to Equation 1, a full fractional design will lead to 1,469,664 ( $3 \times 7 \times 3 \times 4 \times 3 \times 3 \times 3 \times 4 \times 2 \times 3 \times 3 \times 3$ ) simulations, which is non-viable. Therefore, the use of a fractional factorial design is required.

The study has been carried out in two stages: the initial design and the final design. With the initial design the aim is to obtain the simple and double effects that are significant. In the final design only those effects worthy of interest will be studied in detail.

**3.5.1. Initial design** At this stage, only the design factors have been considered. The noise factors and two levels per factor have been left out. The minimum and maximum levels of each factor are shown in Table 1. The cho-

sen design is a  $2^{11-4}$  design, which allows us to study all the simple effects and double interactions between the considered factors. After making the 128 simulations the corresponding ANOVA table shows that only these effects are significant: TROUT, TARB, FCROSS, TABLE, POLIC, NVC, TFLY\*POLIC<sup>3</sup>, TROUT\*TARB, TROUT\*TABLE, TROUT\*POLIC, TROUT\*NVC, TARB\*FCROSS, TARB\*TABLE, FCROSS\*TABLE, TABLE\*POLIC, TABLE\*NVC, POLIC\*NVC. The final design must include at least the factors that are significant by themselves, or those that are significant in a double interaction. The double interactions themselves must be included for further study.

**3.5.2. Final design** The final design can be carried out using the information obtained from the initial design, but now taking all levels into consideration. This design will allow us to obtain the effect of the design and noise parameters over the network performance.

For this purpose it is usual to use one of the orthogonal arrays of Taguchi [9]. However, when the studied factors have a different number of levels, the selected design must be adapted. In this case, an orthogonal array has been used to accommodate those factors with three or less levels. The resulting design has been crossed with the remaining factors to obtain the final design.

The orthogonal array  $L_{81}$  is the smallest one that is able to accommodate the factors with three or less levels, and their interactions. Although the initial design showed that it was not necessary to include MUX and TFLY factors, they have been added without increasing the number of tests. This has been done because a quadratic relation could not be correctly observed on a two level analysis, and in some cases conclude that a significant factor is not significant.

<sup>3</sup> A\*B means a double interaction between A and B.

The 81 combinations of the  $L_{81}$  have been crossed (in a full factorial way) with the four levels of TABLE, the three levels of MESSIZE and the seven levels of DDEST, for a total of 6804 ( $81 \times 4 \times 3 \times 7$ ) simulations. This number could have been reduced through more fractional factorial design, but it allows us to study a large number of double interactions with a high precision (especially interactions between noise and design parameters), being simulated in a reasonable time. Now it is possible to calculate the effect of each design parameter as the average of a great number of data. The effect of a factor with three levels, as the TROUT for example, it is calculated with  $6804/3 = 2268$  observations.

#### 4. Results

The execution of the experiments of the final designs produced a great deal of information. With the results of the 6804 simulations, the results of the combinations that have not been executed can be estimated. Furthermore, many effects and interactions have been extracted. The ANOVA table (Table 2) shows if the simple and double effects are significant or not. Due to lack of space, only the more interesting results will be shown. The more interesting interactions are shown in Figure 2.

- **FCROSS\*TABLE**: With a parallel access when the number of bytes per cross increases, then the throughput also increases. This is because at a certain moment, with a parallel access to the routing table there are more packets ready to cross. This allows a larger number of packets to cross at the same time, taking advantage of a crossbar with a higher capacity.
- **DDEST\*POLIC**: VOQ decreases the negative effect of having a large number of messages whose target is the same hot spots. The problem with hot spots is the congestion of some target hosts and the HOL Blocking. This is because with VOQ and a single switch, the traffic to the different target hosts is completely isolated, and HOL blocking is not a problem.
- **MESSIZE\*TABLE**: The longer the messages, the smaller the difference between the different options on the number of ports of the routing table. This is specially so in the case of having four and infinite ports. This is because at the same injection rate, the number of packages decreases, reducing the need to route a high number of messages at the same time.
- **NHOST\*NVC**: The higher the number of hosts (in this case, ports of the switch), the better the results with a higher number of virtual channels. With a high number of ports, the problem of HOL Blocking increases. With a high number of virtual channels this problem decreases. This is because having a high number of vir-

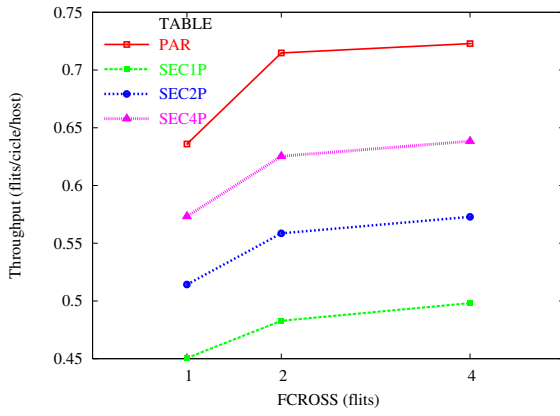
Source	Sum of Squares	Df	Mean Square	F-Ratio	P-Value
<b>MAIN EFFECTS</b>					
A:TPLY	4,0527	2	2,02635	145,13	0,0000
B:TROUT	44,0708	2	22,0354	1578,16	0,0000
C:TARB	0,724057	2	0,362029	25,93	0,0000
D:BCROSS	4,71818	2	2,35909	168,96	0,0000
E:MIX	0,410442	1	0,410442	29,40	0,0000
F:TABLA	37,7228	3	12,5743	900,56	0,0000
G:LONGMEN	50,9163	2	25,4582	1823,29	0,0000
H:DDEST	41,1005	6	6,85008	490,60	0,0000
I:POLIT	99,6682	2	49,8341	3569,08	0,0000
J:NHOST	120,491	2	60,2457	4314,75	0,0000
K:NVC	5,65224	2	2,82612	202,40	0,0000
<b>INTERACTIONS</b>					
AF	1,4558	6	0,242634	17,38	0,0000
AG	0,695939	4	0,173985	12,46	0,0000
AH	0,893614	12	0,0744678	5,33	0,0000
BC	3,34881	4	0,837204	59,96	0,0000
BF	22,8831	6	3,81384	273,14	0,0000
BG	21,553	4	5,38824	385,90	0,0000
BH	2,65394	12	0,221161	15,84	0,0000
BI	11,4654	4	2,86636	205,29	0,0000
BK	1,56129	4	0,390324	27,95	0,0000
CD	1,86744	4	0,466859	33,44	0,0000
CF	0,622631	6	0,103772	7,43	0,0000
CG	1,21209	4	0,303022	21,70	0,0000
CH	0,157258	12	0,0131048	0,94	0,5066
CI	3,38927	4	0,847319	60,68	0,0000
CK	0,153635	4	0,0384088	2,75	0,0266
DE	0,117661	2	0,0588303	4,21	0,0148
DF	0,38225	6	0,0637084	4,56	0,0001
DG	0,206911	4	0,0517278	3,70	0,0051
DH	0,922434	12	0,0768695	5,51	0,0000
DJ	0,383902	4	0,0959755	6,87	0,0000
EF	0,311052	3	0,103684	7,43	0,0001
EG	0,236182	2	0,118091	8,46	0,0002
EH	0,186868	6	0,0311447	2,23	0,0375
EI	3,48338	2	1,74169	124,74	0,0000
EK	0,564855	2	0,282427	20,23	0,0000
FG	19,5663	6	3,26104	233,55	0,0000
FH	2,78497	18	0,15472	11,08	0,0000
FI	7,08623	6	1,18104	84,59	0,0000
FJ	4,64111	6	0,773518	55,40	0,0000
FK	0,595252	6	0,0992087	7,11	0,0000
GH	4,83771	12	0,403143	28,87	0,0000
GI	3,34244	4	0,835609	59,85	0,0000
GJ	3,9199	4	0,979974	70,19	0,0000
GK	0,132665	4	0,0331661	2,38	0,0498
HI	17,7898	12	1,48248	106,17	0,0000
HJ	3,15534	12	0,262945	18,83	0,0000
HK	1,14776	12	0,0956471	6,85	0,0000
IK	3,67866	4	0,919665	65,87	0,0000
JK	3,93143	4	0,982857	70,39	0,0000
RESIDUAL	91,2325	6534	0,0139627		
TOTAL (CORRECTED)	688,196	6803			

All P-ratios are based on the residual mean square error.

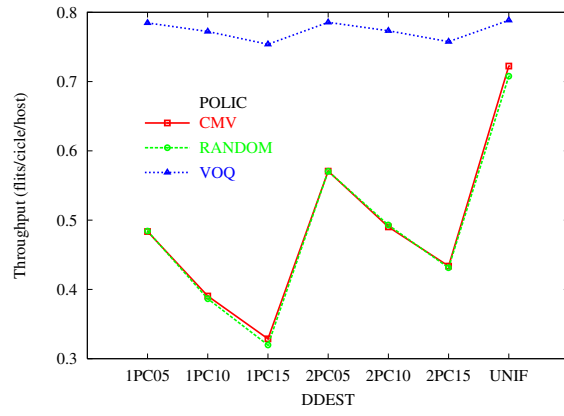
**Table 2. Table Anova for the final design**

tual channels is more important with a high number of ports.

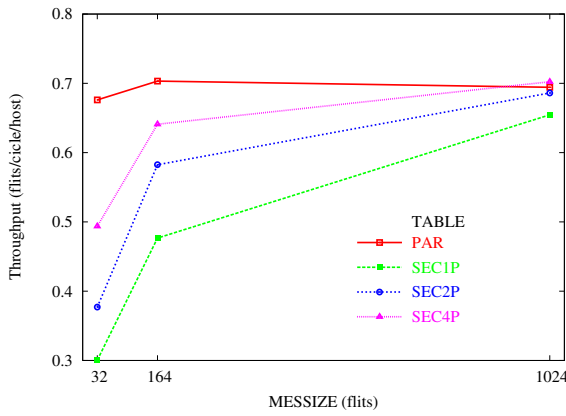
- **TROUT\*MESSIZE**: The longer the messages, the smaller the significance of having a high routing time. This is because with a lower number of messages the routing must be performed a lower number of times.
- **TROUT\*TABLE**: When the routing time increases, the positive effect of increasing the number of ports in the routing table is higher. A high routing time exacerbates the problem of routing table having few ports. This is because more messages have to wait until a port is freed. More messages now arrive at the switch while the routing is in process, these messages also have a longer wait.



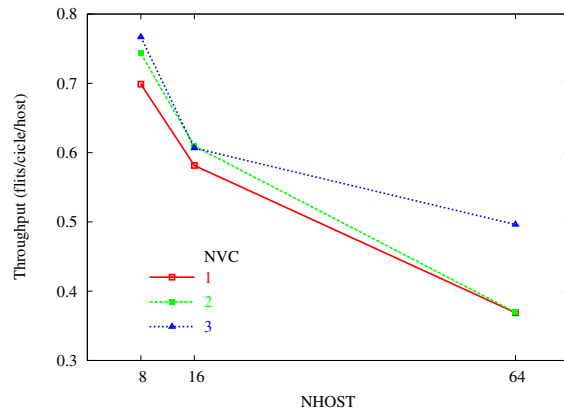
(a) FCROSS\*TABLE



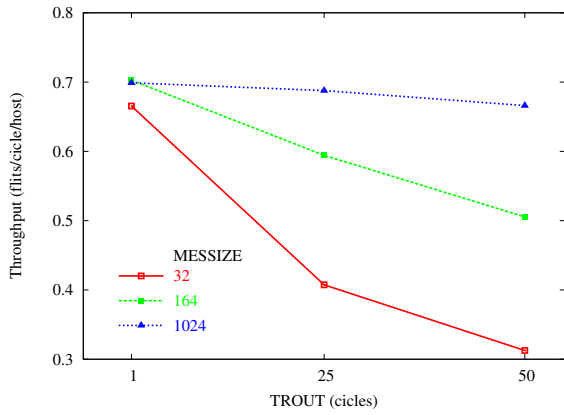
(b) DDEST\*POLIC



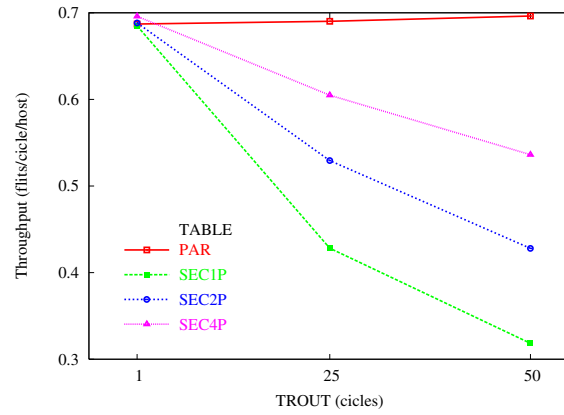
(c) MESSAGE\*TABLE



(d) NHOST\*NVC



(e) TROUT\*MESSAGE



(f) TROUT\*TABLE

**Figure 2. Some double interactions between factors**

## 5. Conclusions and future work

In this paper, the high dependency between the different design parameters in a switch has been shown. This means that the effect of one parameter over the performance of the switch depends on the value of some other parameters. This should be taken into account when carrying out a study of the effect of some concrete design parameter.

With the appropriate design not all the possible combinations of the parameter values require testing. With a small number of simulations a high amount of information can be obtained. In this case, it would be necessary to simulate more than one million cases to consider all the possible combinations. With the technique used only 6804 simulations have been needed for the final design. Another relevant aspect is that it is not necessary to consider a high amount of different values for each parameter. For the quantitative parameters only three values are sufficient to study quadratic effects. This points to the conclusion that a good experimental design is necessary in order to obtain the maximum information with the minimum number of experiments.

The first conclusion is that a study of this kind can produce a very large amount of information that requires further study to obtain concrete results.

As regards to the specific results obtained, it must be said that some of them were totally in line with expectatives both logical reasons and in the light of previous knowledge. However, other results are more difficult to explain and would require additional studies in order to obtain more information.

The methodology used and the characteristics of the results obtained pave the way for further studies. These studies could be more specific and more contained the present study or more ambitious. Thus, some possible future work could focus on the study of some of the obtained results. One of these studies could be centered on the TABLA fac-

tor. The goal of this study could be to determine how many ports would be necessary at the routing table to obtain the advantages of parallel access. Other studies could be variations of this one, as for example, the study of a different set of factors or an evaluation with other performance metrics. Our study has been carried out with a single switch. In this sense, an interesting study could be to focus on a network with multiple switches. This would also allow us to consider other design parameters, as for example the routing algorithm or the topology.

## References

- [1] N. Boden, D. Cohen, and R. Felderman. Myrinet – A Gigabit per second local area network. *IEEE Micro*, pages 29–36, Feb. 1995.
- [2] R. Buyya. *High performance cluster computing*. Prentice Hall, 1999.
- [3] J. Duato, S. Yalamanchili, and L. Ni. *Interconnection networks. An engineering approach*. IEEE Computer Society, 1997.
- [4] A. A.-D. et al. High performance sorting on networks of workstations. In *Proceedings of the ACD SIGMOD*, May 1997.
- [5] InfiniBand Trade Association. *InfiniBand Architecture Specification Volume 1. Release 1.0*, Oct. 2000.
- [6] R. Jain. *The art of computer systems performance analysis: techniques for experimental design, measurement, simulation and modeling*. John Wiley & Sons, Inc., 1991.
- [7] R. Martínez. Studying the effect of the design parameters on the interconnection network performance in NOWs. Master's thesis, Dpto. Informática, Universidad de Castilla-La Mancha, 2003.
- [8] R. Sheifert. *Gigabit Ethernet*. Addison-Wesley, Apr. 1998.
- [9] G. Taguchi. *System of Experimental Design. White Plańis*. New York. Ed. UNIPUB/Kraus International Publications; Dearborn, Michigan, Ed. American Supplier Institute, Inc, 1985.