

Cost/Performance evaluation of Banyan networks for use in ATM switching fabrics

Christos Bouras

*Department of Computer Engineering and Informatics
Patras University, Rio, 265 00 Patras, Greece*

and

*Computer Technology Institute
3 Kolokotroni St., 262 21 Patras, Greece
email: bouras@cti.gr*

Christos Gkantsidis

*College of Computing
Georgia Institute of Technology
801 Atlantic Drive, Atlanta, GA 30332-0280, US
email: gantsich@cc.gatech.edu*

December 5, 1999

Abstract

In this paper, we present a model for computing the cost of implementing Banyan networks. We limit our interest in Banyan networks which are used in ATM switching fabrics and are build in VLSI. The cost is given as a function of the characteristics of the network (ie. length of buffers, speed of links, etc). We are focusing in this area because such networks have been used extensively in the design of large ATM switches, which are the core elements of broadband networks. It is well-known that the cost of implementing must be related to the performance of the network. The choices, that the designer may have, impact both the performance and the cost. We will show when a slight increase in performance implies a great increase in cost, in which case it is not cost effective to build a better switching network, and of course the reverse, when a decrease in cost implies a degradation of the performance of the switch.

Keywords: Banyan, switching fabrics, ATM, VLSI, cost evaluation, performance evaluation.

1 Introduction

The growing need for new services, like video on demand and many others, have increased the demand for new networks that can support a large number of users and a variety of services, which include both traditional bandwidth-hungry data services and others that require a guaranted quality of service (QoS) from the network. One solution to this problem is the development of Broadband Integrated Services Digital Networks (B-ISDN). Because ATM is the transport protocol of choice for this kind of networks, a need has arised to build ATM networks that support many customers and demanding services. The core of an ATM network is the ATM switch and thus the afore-mentioned needs must be fulfilled by this device. Large ATM Switches must be build and this must happen in a cost effective way, meaning that we must compromise the performance with the cost. The core of the ATM Switch, which influence in a dominant way both the performance and the cost, is its switching fabric. In this work, we will concentrate on a special architecture for building switching fabrics which is based on Banyan networks. This kind of networks are very efficient in supporting a large number of customers and have been used extensively in the past for this purpose. Our results, techniques and methodology can be applied easily to other architectures with slight modifications.

In this paper, we try to identify the factors that affect the design of a Banyan network. Some of these factors are the existence or not of buffers, the length of them if they exist, the presence or not of a sorting network that precedes the Banyan, the speedup of the network and many others. The designer can experiment with them to determinate which network suits better the design goals. The literature is rich

with works that examine the impact of these factors in the performance of the network using both analytical techniques and experimental results. We are using some of them to evaluate the impact of some choices on the performance of the Banyan network.

But, except for the performance of the switching fabric there is another critical parameter that affect the design and this is the cost. A choice which can increase the performance slightly, but will have a significant impact on the cost is not wise, except in rare cases where the only goal of the design is the performance. On the other hand, a small increase in cost which results in a great advancement in performance should be made. Assuming that the switching fabric will be implemented in VLSI, a major parameter of the cost is the area of the chip. This area affects the probability of producing a working chip because the probability of a fault is constant in a unit area and depends only on the technology being used. Thus, a small chip has greater probability to be produced without a fault. Increasing the area, by adding for example buffers in order to improve the performance of the switch, result in an increase in the probability of producing a faulty chip. Chip area is not the only factor that affect cost; another one is the speed up factor which shows how faster will be the links of the swithing fabric than the input and output links. In this work we are trying to relate the cost with these parameters. Having found a way to express the cost with the parameters of the network, we will compare the cost and the performance and suggest a method to find the optimal cost performance design.

Many papers found in the literature, like [Copp93], [Fran81] and [Bour98], are the basis of this work and many of our results depend on these papers. Although the key points of them are mentioned in this paper, the interesting reader must consult them for full proofs and explanations.

2 The model

2.1 Banyan networks

Banyan networks belong to the class of Multistage Interconnection Networks (MINs). They were defined in [Goke73] and have the property that there is exactly one path from any input to any output. A simple example of a 8×8 banyan interconnection network (ie. which has 8 inputs and 8 outputs) is drawn in figure 1. This network consists of nodes and links. Every node is a 2×2 switch, which can receive packets at each of its two input ports and send them through each of its two output ports. Generally, switches may have k input ports and k output ports¹ ($k \times k$ switches). Switches are grouped in stages. This means that every switch in stage j can be connected to switches in stages $j - 1$ and $j + 1$. Of partucial interest are networks in which all switches are identicall, except perhaps the switches in the first and last stages. The other building element of a MIN is links. Links are used to connect switches of successive stages.

¹we will concentrate on switches that have the same number of inputs as outputs

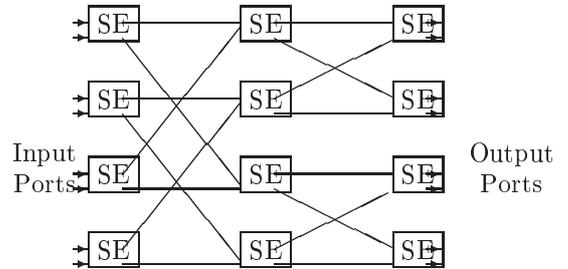


Figure 1: An 8×8 Banyan network

So far, we have made some assumptions about the switches. Now we will concentrate on the network. We will assume that the problem is to interconnect N inputs to N outputs, thus we want to construct an $N \times N$ MIN which has a banyan topology. Using $k \times k$ switches, we need a network with $\log_k N$ stages and $\frac{N}{k} \log_k N$ switches. We will also assume that we want to transport packets from inputs to outputs without preestablishing a path; our transport network will be connectionless.

Banyan networks have many advantages. First of all they require only $\frac{N}{k} \log_k N$ switches and $N \log_k N$ connections, as opposed to the crossbar network which requires $O(N^2)$ switches and links. Also, they have the property of self-routing. This means that every switch that accepts a packet in one of its inputs can decide in which of its output to forward this packet depending only on the destination address of the packet. For example, in the Delta-2 networks, which are a subclass of Banyan networks consisting of 2×2 switches, every switch of stage j can decide in which output port to forward a packet based on the j bit of the destination address. If this bit is 0 then the packet is forwarded to the "upper" output port, and if it is 1, it is forwarded to the "lower" output port. Delta networks were defined and described in [Pate79], [Pate81]. Another advantage of Banyan networks is their regularity and interconnection patterns which are very attractive for VLSI implementation ([Ahma88]).

But, these advantages come with a cost. It is well known that Banyan networks are blocking. This means that it is possibly for two or more packets to content for the same link somewhere in the network. Only one of them will win this contention and will be transmitted. The others must be buffered and try again in a later time (we will assume that the network must not discard packets if there are enough buffers). So, a choice that a designer may have is where to place the buffers. If he or she places them in the input ports, which means that a packet stays there until it is finally transported to an output port, then due to head-of-line blocking² the throughput of the network will be much lower than $2 - \sqrt{2} \approx 0.586$. We will not prove this limit or

²the packet in the head of the queue will prevent the other packets from reaching an output port even if they can find a path in the network

discuss the assumptions, which are close to the ones we have made so far; the interesting reader may refer to [Pryc95] or [Karo87]. Another option is to place buffers in the switches. A lot of research has been done concerning this matter and a lot of results, both theoretical and from simulations, have appeared (see for example [Bour98], [Dias81], [Krus83]).

The designer has three other options to increase the performance of the switching fabric. The first is to increase the speed of the internal links by a factor of s . More info about this can be found in [OIE89]. Also, he can use d -dilated or d -replicated Banyan networks which were introduced in [Krus83]. The third option is to use Batcher-Banyan networks, or more generally sorting banyan networks, which were introduced in [Batc68] (see also [Ahma88]). The most interesting characteristic of them is that they are internally non-blocking.

2.2 Factors that affect the cost of the networks

From the description of Banyan network is easily derived that the number of required switching elements (SE for short) is $O(N \cdot \log N)$ and that the number of stages is $O(\log N)$, where N is the number of input/output ports in a $N \times N$ transport network. This means that the total chip area that must be used is $O(N \cdot \log N)$ when the design is based on SSI technology. But when VLSI is used this doesn't happen as Franklin points out in [Fran81].

Traditional techniques for finding the total area of the chip take for granted that the cost of interconnecting subsequent stages is negligible and this means that they assume that the total area depends only on the number of switching elements. Let's assume that each line, which connects a switching element in one stage to another element in the following stage, has a width of l units of length. Let's also assume that the gap between adjacent lines must be l units of length. Thus, if w parallel lines are needed to connect two SEs (path of length w), then at least $2lw$ units of length are required for this connection. The bad thing about Banyan networks is that connection paths must be mixed up in some stages. Take for example the network in figure 1, where it is clear that the paths between the first and second stages have more points in common than the paths between the second and the third stage. The designer must find a way to avoid these common points given that in VLSI the paths consist of horizontal and vertical segments (ie. there are no diagonal connections as in figure 1). As indicated in figure 2, which is the same interconnection network as in figure 1 with the exception that the above constraint about communications paths has been taken into account, in order to construct the paths which connect two subsequent stages when there are many points in common an increase in the horizontal distance of the stages must be made. In a network with more inputs and outputs which is constructed using the pattern of figure 1, the distance between subsequent stages would be a decreasing function. What is most interesting is that the total area which is used to connect the SEs of states, which is the summation of the products of the vertical times the

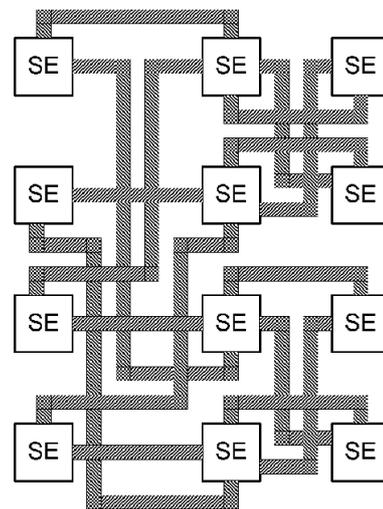


Figure 2: An 8×8 Banyan network implemented in VLSI

horizontal distances, is $O(N^2)$. The interesting reader is referred to [Fran81] for a full proof and a deeper explanation of the above reasoning. Something that we must point out, is that in Franklin's work it was assumed that the SEs had no buffers. If we put buffers in the switching elements then the area of the Banyan network will be dominated by the total area of the SEs.

A large part of the previous discussion has been devoted to methods which will help us to estimate the required area for the Banyan transport network. We haven't said enough for the importance of this estimation and its relation to the cost function, which is our ultimate goal. It is well known that when producing a number of chips, using any available technique, not all of them will work properly. In fact the vast majority of them will be faulty. The probability of obtaining a working component of unit area in a given technology is called yield and we will denote this by r . Thus if our chip, which is a Banyan network, has area equal to A units, then the expected number of correct chips will be one every r^{-A} . Now, we must try to relate this probability and the area A with the cost of producing the chip. Generally this is a very difficult task. To simplify it, we will adopt the approximation, which is also made in [Copp93], that the cost depends on $A \cdot r^{-A}$.

Except from the area of the chip, we will assume that the cost depends also on the speedup factor. It is evident that this is true for each link; stated in a different way the cost of each link depends on the speedup factor. Also, the cost of each switching element depends on this, because increasing the operational speed of the links means that we must increase the speed of the SE. Another assumption that we will make is that the total cost depends on the cost of each link and of each switching element.

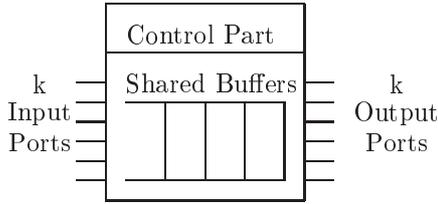


Figure 3: An abstract structure of a SE

3 Analytic results

3.1 Evaluation of the function of the cost

In this section we try to evaluate the cost of the networks discussed in the previous section. First, we turn our attention to the basic building block of the Banyan network, which is the switching element. Every SE must have the structure of the figure 3. It has k inputs and k output ports. Also, there are at least k buffers in the SE, which hold the packets temporarily before they are transmitted in the next stage. In a non-buffered network each SE will have exactly k buffers. In buffered networks, it is usually assumed that each output port is associated with B buffers. Thus there are $k \cdot B$ buffers in the SE. This doesn't happen in practice. Usually, fewer than $k \cdot B$ buffers are used which are shared between all output ports. By multiplexing all output streams (packets for the same output port) the number of buffers can be reduced without affecting the performance, in terms of cell loss probability or average delay. The penalty for this is an increase in the complexity of the SE and the use of faster, by a factor of at least k , memory. The ratio of the number of buffers when shared memory is not used, which is $k \cdot B$, and the number of buffers in the other case when the same performance is achieved, is called multiplexing gain and depends on the number of inputs and the traffic characteristics. We will make the assumption that traffic is uniformly distributed and thus this ratio will depend only on k . We expect that as k increases then this ratio will also increase because better statistical behavior would be achieved³. We denote the reciprocal of this ratio as $m(k)$ and thus the total number of buffers that will be used is $k \cdot m(k) \cdot B$. The size of the area which will be used for the buffers, A_B , depends on the number of them and this means that:

$$A_B = O(k \cdot m(k) \cdot B) \quad (1)$$

Function $m(k)$ is monotonically decreasing and takes values in $(0, 1]$.

Eq.1 gives an approximation of the area of the SE which will be used for the buffers. In order to find the total area of the SE, we must approximate the area of the control part

N	Number of input and output ports of the switching fabric.
k	Number of input and output ports of the switching elements.
s	Speedup factor.
d	d -replicated or d -dilated networks.
r	Yield. The probability of producing a working chip of unit area.
B	Number of buffers per output port.
$m(k)$	Multiplexing factor of input streams in $k \times k$ SEs.
$m'(k)$	Same as $m(k)$. They are used interchangeably.
A	Area of something.
A_B	Area used for buffers.
A_{CP}	Area used for control part.
A_{SE}	Area of a single SE when buffers are used.
$A_{total SE}$	Total area needed for SEs in a buffered Banyan network size N .
C_{BB}	Total cost of buffered banyan networks.
C_{DB}	Total cost of d -dilated or d -replicated banyan networks.
α_B	Constant depending on the implementation of buffers.
α_{CP}	Constant depending on the implementation of the control part.
α_{SE}	Constant depending on the implementation of the SE.
α_I, α'_I and α''_I	Constants depending on the implementation of the links.
k_c, k'_c and k''_c	Constants depending of the technology used to build the chip.

Table 1: List of symbols and their meaning

(see fig.3). Clearly, this depends on the number of input and output ports. Also, we expect that if many buffers are used then this area will be much smaller than the area of the buffers. On the other hand, when no buffers are used then we expect that the area of the control part is comparable, if not bigger, to the area of the buffers. Thus, if we assume that α_B and α_{CP} are constants which depend on the implementation, then the total area of the SE is⁴:

$$A_{SE} = A_B + A_{CP} = \alpha_B \cdot k \cdot m(k) \cdot B + \alpha_{CP} \cdot k$$

An approximation which has been made in [Copp93] and which we will use is that the right part of the previous equation is $O(k \cdot m'(k) \cdot B)$. In this way, we use the new function $m'(k)$ to hide the area of the control part. Thus, the total area of the switching element is:

$$A_{SE} = O(k \cdot m'(k) \cdot B) \quad (2)$$

By introducing the constant α_{SE} we can rewrite the previous equation as:

$$A_{SE} = \alpha_{SE} \cdot k \cdot m'(k) \cdot B$$

From now on, $m(k)$ will refer to $m'(k)$.

The number of SEs of the Banyan network is $\frac{N}{k} \cdot \log_k N$. The area they consume is:

$$A_{total SE} = O\left(\frac{N}{k} \log_k N \cdot k \cdot m(k) \cdot B\right) \quad (3)$$

Observe that the above equation gives only the necessary area to implement the functions of the SE (buffering and control). In practice, it is possible to use more area because

⁴See table 1 for a listing of all parameters and constants used in this work.

³a fact that is derived from the central limit theorem

we must take into account that only a limited number of positions can be used to place the SEs and that the length of the sides of the SEs must be sufficiently large to place the input and output ports. This later approach has been used in [Fran81], where it is proved that the total area needed for the SEs is $O(N^2)$. In contrast, eq.3 says that this area is $O(N \cdot \log N)$. In our analysis we assume that the area of the buffers, when they are used, dominates the total area of the SE. The remaining area, which is asymptotically $O(N^2)$, is included in the area needed for the connections, which is given below.

The remaining area of the chip is used for the links. Even though the necessary number of links is $N \cdot \log_k N$, the area they use is $O(N^2)$ (see [Fran81]). The cost of the links depends also on the speed up factor s . Thus, we can approximate this cost with the function $\alpha_I \cdot N^2 \cdot s$, where α_I is a constant which depends on technology being used to implement the chip. The extra cost which is needed for the reasons discussed above can be included by increasing this constant.

Now, we can give the total cost of a Banyan network which uses buffers. This cost depends on the area used for the SEs and the links, on the speedup factor and on yield and it is:

$$C_{BB} = \frac{k_c \cdot s \cdot N \cdot (\alpha_{SE} \log_k N \cdot m(k) \cdot B + \alpha_I \cdot N)}{r^{-N \cdot (\alpha_{SE} \cdot \log_k N \cdot m(k) \cdot B + \alpha_I \cdot N)}} \quad (4)$$

where k_c is a constant which depends on the implementation and the technology used. Constants α_{SE} and α_I must be selected in such way that they accurately give the areas used by SEs and links respectively. Thus, in a buffered network where the area of the SE dominates the area of the chip, α_{SE} must be larger than α_I so that $N \cdot \log_k N \cdot m(k) \cdot B$ is much bigger than N^2 in the previous equation for all practical values of N . When only k buffers per SE are used (no extra buffers, unbuffered network) then the area of the links dominates the area of the chip and this means that α_I is bigger than α_{SE} .

The cases of d -replicated and d -replicated unbuffered Banyan networks are analogous to the basic network. d -Replicated networks consist of d copies of simple Banyan networks. Assuming that the area of control logic is negligible, the cost of these networks is

$$C_{DB} = k'_c \cdot d \cdot \alpha'_I \cdot N^2 \cdot r^{-d \cdot \alpha'_I \cdot N^2} \quad (5)$$

The case of d -dilated networks is somewhat more complicated. The total area of the links, which are used in this case, increases by a factor of d . Thus, the above equation holds in this case, too.

Finally, we examine the Sort-Banyan networks. These networks consist of two subnetworks which are called sorting and banyan. Usually the sorting subnetwork is a Batcher bitonic. In this case the SEs of this subnetwork are 2×2 switches. In practice all SEs are identical and thus the total number of SEs required is

$$\text{Number of SEs} = \frac{N}{4} \cdot \log_2^2 N + \frac{3 \cdot N}{4} \cdot \log_2 N$$

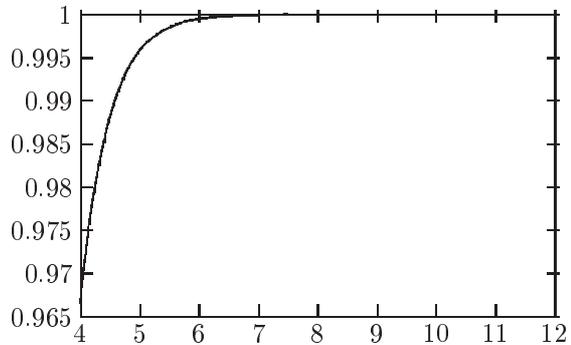


Figure 4: **Gain in cost versus number of ports of SE**

. As in the case of simple Banyan networks, the total area is $O(N^2)$. The main difference here is that the chip area needed for this network is two times the area of the simple Banyan, because two subnetworks are used; one of them is Banyan and the other has a similar structure. Doubling the area doesn't mean that the cost will also be double. It would be much more because it increases exponentially with area.

3.2 Some observations on the functions

Now, let's look how the functions derived in the previous section can help the designer make some decisions which will reduce the cost of implementing the network.

First of all, a choice must be made about the size of the SEs, ie. the value of k . Increasing k means that fewer buffer per output port will be needed because of statistical multiplexing. But, k can't exceed a maximum value, which depends on the technology used. These two facts can lead to conclusion that k must take its maximum value. Using eq.4, we can plot the gain in cost of using $k \times k$ switches (ie. $\frac{C_U - C_B}{C_U}$ where C_U is the cost of the switching element when no multiplexing is used and C_B is the cost in the case of shared buffers) versus number of ports. Even though this plot depends on many factors, a general sketch of it can be found in fig.4. For this sketch, we have used an approximation of $m(k)$ found in [Copp93]⁵; but generally for every valid $m(k)$ the function have the same shape. A property of this function is that it is increasing very rapidly and its limit is one. Also it is easily observed that there is a value of k before of which the rate of increasing is very rapid and after of which this rate is very slow. Thus, the best choice of k , which reduce the cost and keeps the implementation simple, is this value. For example, for $r = 0.3$ and for several values of N and B the optimum is close to 8, meaning that 8×8 switches must be used.

Some other design parameters, such as length of buffers per port B , d in replicated or dilated networks and N , have the reverse effect in the cost. They increase exponentially the cost of building the network.

⁵ $m(k) = \frac{0.35 \cdot k + 2.9}{k + 1.5}$

4 Cost/performance evaluation

In order to study the performance of the switching fabrics, we must define the measures which will be used to characterize their performance. One such measure is the probability of a packet being dropped in the network because of buffer overflows. This is particularly important to the end user whose applications depend on the reliable transportation of information. If this probability is large then lost packets will result in a degradation of the quality of applications. Also, this metric is important to the network, because lost packets can trigger retransmissions which can result in congestion. To avoid this problem, large buffers must be used. Another measure of interest is the total delay which is needed for the packet to traverse the switching fabric. Because we have assumed ATM networks, in which the length of packets is constant, this measure depends only on the number of packets waiting in queues and on the frequency of operation of the switching fabric. In this paper we will concentrate on the expected time to traverse the fabric and thus we will use the mean number of packets waiting in queues. The last measure we will use is the throughput of the switching fabric. This is important to the holder of the network who want to forward as many packets as possible, or, saying it in other words, to make best use of the network.

An interesting thing to study is the effect of using $k \times k$ switches with large k in the probability of dropping a packet. In [Bour98] it was conjectured that using small k 's is preferable when it comes to this probability. Using the algorithm found in that paper we have been able to construct the graphical representation of fig.5. This figure represents the mean number of packets lost per queue in the first stage of a buffered Banyan network versus k , and it is also valid not only for other stages but for unbuffered networks, too. On the other hand, in the previous section we have found that increasing k has the advantage of using fewer buffers per output port and thus decreasing the total area needed for the SE. In fig.6 we can see the ratio of the mean number of packets lost per output queue and the gain of using multiplexing (found in the previous section) versus k . In this figure smaller values are better. We can easily observe that increasing k increases slightly this ratio. Thus, it is preferable to use small k .

Another interesting parameter of the buffered-Banyan networks is the number of buffers in each output queue. Increasing this number gives a better (ie. smaller) probability of dropping a packet but also results in an increase in the area needed for the chip and thus in the cost. Fig.7 gives this probability versus the size of the buffers for some stages of a buffered Banyan switching fabric. This probability is decreasing very fast but, on the other hand, the cost is increasing exponentially, as we have proved in the previous section. The product of this probability and the cost is given in fig.8 and once again small values are better. It is easily observed that this product is increasing quite fast (exponentially) and this means that adding few buffers beyond a point increases dramatically the product and thus it is not a wise choice to increase the buffers.

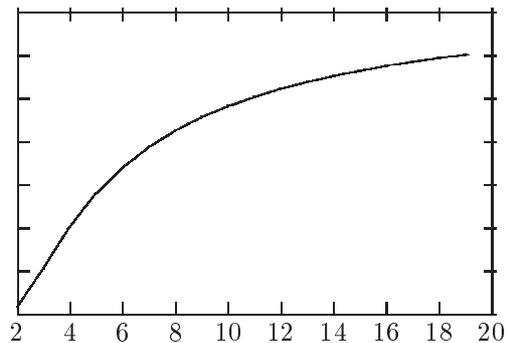


Figure 5: Mean number of packets being lost per queue in the first stage of a buffered Banyan network versus number of ports of SE

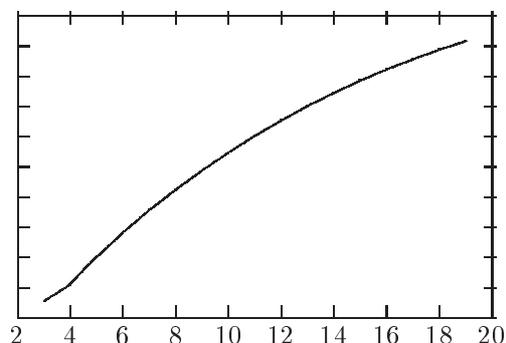


Figure 6: Mean number of packets lost per queue over gain of multiplexing versus k

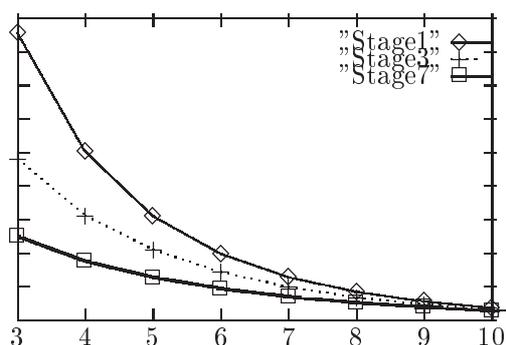


Figure 7: Mean number of packets being lost per queue of a buffered Banyan network versus the size of the queue

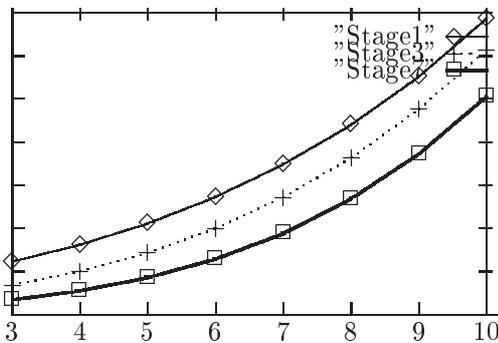


Figure 8: Mean number of packets lost per output queue times the cost versus k for the first, the third and the seventh stages

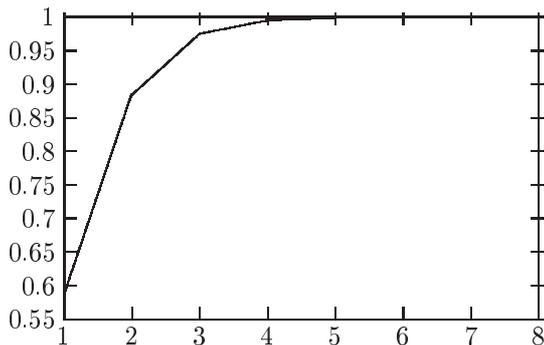


Figure 9: Throughput versus speed-up factor s in Batcher-Banyan switching fabrics

Another interesting parameter is the speedup factor s . The effect of speedup has been extensively studied in the literature for the case of non-blocking switches; Batcher-Banyan switches fall in this category. In [Chen91] and [OIE89] was shown that the throughput of the switching fabric increases as indicated in fig.9. We can observe that when s is small enough the increase in throughput is tremendous. On the other hand choosing a speedup factor of $s + 1$ instead of s will result in a small increase when s is large (practically zero increase for $s > 9$). The price it is paid to increase the throughput is a linear increase in cost. The value of s impacts not only the throughput, but also other performance measures such as the average system delay (ie. the average number of queued packets) and the probability of a packet drop due to buffer overflow. Increasing s will result in fewer queued packets and smaller loss probability (see [Chen91] and [OIE89]). Thus, the designer must choose to use a speed up factor between 3 and 6, whenever this is possibly, because in this way the increase in the performance will overwhelm the increase in cost. Even though the above results apply to nonblocking switching fabrics, such as Batcher-Banyan, we expect that the general conclusions will also apply to buffered Banyan networks.

Using the same techniques for the case of d -replicated and d -diluted networks, we can prove that it is preferable to use many parallel networks or links from a cost/performance point of view (assuming that the performance is character-

ized by the probability of packet loss).

So far, we have studied the impact that the parameters of the switching fabric may have both to its cost and performance. Using various cost/performance measures we have seen that a change in a parameter which results in an increase in performance, in most cases, results in a bigger increase in cost. Also, the cost in these cases increases exponentially, which means that there is a value of the parameter in question which is a threshold. Moving from a point, which is bigger than this threshold, to another point close to it, results in a great increase in cost. On the other hand, moving between points, which are smaller than this threshold result in relatively small increases.

5 Conclusions — Future work

In this paper we have concentrated in the analysis of various types and extensions of Banyan network both from performance and cost of implementing in a single chip point of view. We have found that using small SEs is better from a cost/performance perspective, that replicated networks can decrease substantially the probability of packet loss without increasing the cost too much and that usually an increase in performance result in a bigger increase in cost. Also, we have found that in most cases it is possible to increase the performance without increasing the cost by many factors, and this is happening when the value of the parameter, which increases the performance, is below a threshold. If the designer keeps this rules, then the design of low-cost switching elements and generally ATM Switches can be happen without reducing the performance too much.

References

- [Goke73] L.R.Goke and G.J.Lipovski. *Banyan Networks for partitioning multiprocessor systems*. in Proc. 1st Annu. Int. Symp. Comput. Architectures, Dec. 1973, pp.21–28
- [Pate79] J.H.Patel. *Processor-memory interconnections for multiprocessors*. in Proc. 6th Annu. Int. Symp. Comput. Architecture, Apr. 1979, pp. 168–177
- [Pate81] J.H.Patel. *Performance of processor-memory interconnections for multiprocessors*. IEEE Trans. Comput., vol. C-30, pp. 771-780, Oct. 1981
- [Ahma88] Hamid Ahmadi and Wolfgang E. Denzel. *A Survey of Modern High-Performance Switching Techniques*. IEEE J. Selected Areas Commun., vol. 7, no. 7, pp. 1091-1103, Sep. 1989
- [Karo87] Mark J. Karol, Michael G. Hluchyj and Samuel P. Morgan. *Input Versus Output Queueing on a Space-Division Packet Switch*. IEEE Trans. Commun., vol. COM-35, no. 12, pp. 1347–1356, Dec. 1987

- [Pryc95] Martin de Prycker. *Asynchronous Transfer Mode: Solution for Broadband ISDN*. Third Edition, Prentice-Hall, 1995
- [Bour98] C.Bouras, J.Garofalakis, P.Spirakis, V.Triantafillou. *An analytical performance model for multistage interconnection networks with finite, infinite and zero length buffers*. Performance Evaluation, vol. 34, pp. 169-182, 1998
- [Dias81] Daniel M. Dias and Robert Jump. *Analysis and Simulation of Buffered Delta Networks*. IEEE Trans. Comput., vol C-30, no. 4, pp.273-282, Apr. 1981
- [Krus83] Clyde P. Kruskal and Marc Snir. *The Performance of Multistage Interconnection Networks for Multiprocessors*. IEEE Trans. Comput., vol. C-32, no. 12, pp. 1091-1098, Dec. 1983
- [OIE89] Juji Oie, Masayuki Murata, Koji Kubota and Hideo Miyahara. *Effect of speedup in nonblocking packet switch*. IEEE International Conference on Communications '89, pp. 410-414, June 1989
- [Batc68] K. E. Batcher. *Sorting networks and their applications*. in Proc. Spring Joint Comput. Conf., AFIPS, pp. 307-314, 1968
- [Fran81] Mark A. Franklin *VLSI Performance Comparison of Banyan and Crossbar Communications Networks*. IEEE Trans. on Computers, vol. C-30, no. 4, Apr. 1981
- [Copp93] Paolo Coppo, Matteo D'Ambrosio and Riccardo Melen *Optimal Cost/Performance Design of ATM Switches*. IEEE/ACM Transactions on Networking, vol. 1, no. 5, Oct. 1993
- [Chen91] Jeane S.-C. Chen, Thomas E. Stern. *Throughput Analysis, Optimal Buffer Allocation, and Traffic Imbalance Study of a Generic Nonblocking Packet Switch*. IEEE J. Selected Areas Commun., vol. 9, no. 3, pp.439-449, Apr. 1991
- [Lea90] Chin-Tau Lea. *Multi-Log₂N Networks and Their Applications in High-Speed Electronic and Photonic Switching Systems*. IEEE Trans. Commun., vol. 38, no. 10, pp. 1740-1749, Oct. 1990