

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCS-80-02

1980-04-01

VLSI Performance Comparison of Banyan and Crossbar Communications Networks

Mark A. Franklin

The performance characteristics of banyan and crossbar communications networks are compared in a VLSI environment where it is assumed that the entire network resides on a single VLSI chip. A high level model of the space (area) and time (delay) requirements for these networks is developed and relative performance comparisons are made based on space-time product measure. The results differed significantly from those obtained with more traditional analyses which are usually based on switch aggregate comparisons. These analyses usually lead to the banyan as being the preferable network due to the $N \log_2 N$ (for banyan) versus N^2 (for... [Read complete abstract on page 2.](#)

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research

Recommended Citation

Franklin, Mark A., "VLSI Performance Comparison of Banyan and Crossbar Communications Networks" Report Number: WUCS-80-02 (1980). *All Computer Science and Engineering Research*. https://openscholarship.wustl.edu/cse_research/880

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

VLSI Performance Comparison of Banyan and Crossbar Communications Networks

Mark A. Franklin

Complete Abstract:

The performance characteristics of banyan and crossbar communications networks are compared in a VLSI environment where it is assumed that the entire network resides on a single VLSI chip. A high level model of the space (area) and time (delay) requirements for these networks is developed and relative performance comparisons are made based on space-time product measure. The results differed significantly from those obtained with more traditional analyses which are usually based on switch aggregate comparisons. These analyses usually lead to the banyan as being the preferable network due to the $N \log_2 N$ (for banyan) versus N^2 (for crossbar) switch growth. The conclusion of this path is that using a space-time product performance measure, and roughly current technology parameters. Under these conditions the crossbar is probably preferable due to its greater conceptual and implementation simplicity.

**VLSI Performance Comparison of Banyan
and Crossbar Communications Networks**

Mark A. Franklin

WUCS-80-02

April 1980

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899**

*Presented at the Workshop on Interconnection Networks for Parallel and
Distributed Processing, Purdue University, April 1980.*

*This work has been supported in part by NSF Grant MCS 78-20731 and
NCHSR Grant HS03792.*

VLSI PERFORMANCE COMPARISON OF BANYAN AND CROSSBAR COMMUNICATIONS NETWORKS*

Mark A. Franklin
 Departments of Electrical Engineering
 and Computer Science
 Washington University
 St. Louis, Missouri 63130

Abstract -- The performance characteristics of banyan and crossbar communications networks are compared in a VLSI environment where it is assumed that the entire network resides on a single VLSI chip. A high level model of the space (area) and time (delay) requirements for these networks is developed and relative performance comparisons are made based on a space-time product measure. The results differ significantly from those obtained with more traditional analyses which are usually based on switch aggregate comparisons. These analyses usually lead to the banyan as being the preferable network due to the $N \log_2 N$ (for banyan) versus N^2 (for crossbar) switch growth. The conclusion of this paper is that using a space-time product performance measure, and roughly current technology parameters, there is little difference between the two networks. Under these conditions the crossbar is probably preferable due to its greater conceptual and implementation simplicity.

1.0 Introduction

In recent years there has been increasing interest in tightly coupled multiprocessor systems. This has been due both to the enhanced performance possibilities for such systems, and the steady decrease in hardware costs associated with these systems. One central issue in their design relates to the network over which the multiple processors communicate. Clearly as the number of processors increases, the characteristics of this communications network can become critical to overall system performance and cost. This has led to a wide variety of studies aimed at characterizing and quantifying the performance of such networks (1-12). This paper focuses on two network types, the crossbar (13,14,15), and the banyan (16), and compares their performance characteristics in a VLSI (17) environment where it is assumed that the entire communications network resides on a single VLSI chip. Both networks are also assumed to operate in a circuit switched mode; that is an entire designated path through the network is held during transmission of information between input and output ports. This path is held for the length of the transaction and then released.

The comparison is based on a performance measure which includes both chip implementation area required for a network, and the time delay imposed by the network. Thus both space and time are included in the performance criterion. Note that previous work on network evaluation (12) considered bandwidth and number of switches in the network as the principal measures of interest. This makes sense in an environment where wires and connection paths are of negligible cost compared to switching element costs. The situation changes, however, when network and switch implementation takes place on a single chip. In such an environment connection paths may use substantial amounts of the chip area thus reducing the area available to the switch elements themselves, and thereby reducing the size of a network that can be implemented on a given chip. The time delay associated with the paths may also contribute to the overall delay in a nonnegligible manner.

The remainder of the paper explores these ideas more fully. The second section reviews the general characteristics of the crossbar and banyan networks. The third section presents a VLSI implementation model. Expressions for the chip area required by each network are developed as a function of the network size and other key parameters. The fourth section considers the time delay properties of the networks and additional expressions are presented for network delay, again as a function of network size and certain key parameters. An overall space-time performance measure is derived in section five and performance curves presented and discussed. The final section contains a summary and conclusions of the paper.

The principal contributions of the paper relate first to the development of a VLSI oriented model and methodology for comparing communications networks and second, to some particular conclusions associated with a banyan/crossbar network comparison. With regard to the latter, the most interesting conclusion is that using a space-time product performance measure the differences between banyan and crossbar networks are significantly smaller than predicted by traditional analysis, usually based on simple switch count considerations. Given that the crossbar is simpler both from conceptual and design viewpoints, in the VLSI environment it is the preferable network.

* This work has been supported in part by NSF Grant MCS 78-20731 and NCHSR Grant HS03792.

2.0 Crossbar and Banyan Networks: General Characteristics

Figures 1 and 2 show the overall structure of eight input/eight output crossbar and banyan networks respectively. The switch connection patterns permitted in the networks are also illustrated. The figures shown represent a typical multiprocessor configuration of the B-5500 (18) or C. MMP (19) variety, with the P and M notation standing for Processor and Memory subsystems. Routing of messages through the network can be performed on the basis of an address associated with each output node. Certain key properties of these networks can be identified:

1. Full Interconnection Capability: Both networks permit any single input/output connection to be made by placing the appropriate switches in the proper positions.

2. Blocking/Non-Blocking Capabilities: In the crossbar network, as long as each input port addresses a unique output port, no blocking occurs in the network. That is, all messages can be routed through the network simultaneously. In the banyan network this is not the case, and under certain conditions messages going to different output ports will require use of the same path between two switches. Since under the assumed protocol only one message can hold a given path during message transmission, blocking will occur. This blocking reduces the bandwidth of the banyan network and introduces added delays in system operation. Figure 3 illustrates the probability of message blocking as a function of network size N. The network size is defined as the number of inputs (or outputs) in an N input/N output (square) network. It is assumed that all the messages enter the network at the same time, address unique output ports and randomly win path access in the case of path contention. The results were obtained using simulation methods and demonstrate how the blocking probability increases with the network size.

3. Local Routing Capability: In both networks routing of messages through the system can be done on a local basis where each switch within the network controls its position (see Figures 1B and 2B) based on the destination address of the message to be routed. This had been considered elsewhere (15) and will not be pursued here. It is assumed that the networks to be analyzed use local control procedures.

4. Component Costs (A'): In Small Scale Integrated (SSI) network implementations it is usually assumed that cost is directly proportional to the number of switches required by each network. For square networks of size N the cost for a crossbar network (A'_{CB}) varies as N^2 while the cost for a banyan network (A'_{BA}) varies as $(N/2)\log_2 N$.

5. Delay Costs (D'): Given that a path through a network is available, traditional SSI based analysis assumes that delay through the network is primarily proportional to the average number of switches through which the message passes. If output ports are uniformly addressed then the average number of switches a message will pass

through in a crossbar network (D'_{CB}) of size N is N, while for a banyan network this average number is $\log_2 N$. In the banyan case however, there is

the possibility that the path will not be available (i.e., the message will be blocked) and this must be accounted for in determining the average message delay (D'_{BA}). Assume messages are of equal

length and begin and end simultaneously. Also assume a saturated system where messages are always available to be sent at each input port, and a message retry protocol where blocked messages reenter the system again with the next message batch. Under these conditions, if P_N is the probability of a message being blocked in a network of size N, then the average delay through a banyan network is given by:

$$D'_{BA} = \log_2 N(1-P_N) + 2\log_2 N(1-P_N)P_N + 3\log_2 N(1-P_N)P_N^2 + \dots$$

which can be reduced to:

$$D'_{BA} = (\log_2 N)/(1-P_N) \quad [1]$$

6. Performance Measure: One simple overall performance measure for a network is the product of component and average delay costs. For the crossbar and banyan cases these are given below:

$$C'_{CB} = A'_{CB} \cdot D'_{CB} = N^3 \quad [2]$$

$$C'_{BA} = A'_{BA} \cdot D'_{BA} = N(\log_2 N)^2 / 2(1-P_N) \quad [3]$$

For comparison purposes, if we assume that the individual switch complexity for banyan and crossbar implementations are about the same, then it is convenient to consider the ratio C'_{BA}/C'_{CB} .

$$C'_R = C'_{BA}/C'_{CB} = \frac{1}{2(1-P_N)} \left[\frac{\log_2 N}{N} \right]^2 \quad [4]$$

The remainder of this paper considers how this performance measure changes when one moves from an SSI environment to a VLSI environment where the entire communications network is implemented on a single chip. In particular, one must now consider the required chip area for a given network as an important performance component rather than number of switches per se. Furthermore, there must be some consideration of area related parameters such as network layout and topology. These parameters have generally been of low importance in traditional SSI based network comparisons. In terms of network delay, added consideration must now be given to the connection path contributions to overall delay. Finally, the close relationship between layout, area, and delay must be quantified.

3.0 A VLSI Area Model

In order to develop a VLSI area model without performing detailed network design and layout, certain assumptions must be made and general parameters defined. Some of these are design related, while others are technology related. We begin

with those which are design related and apply to both banyan and crossbar networks. The objective initially is to develop equations for the area taken by a single switch.

Consider first the geometric shape of the individual switches which make up the network. While the actual shape of a switch on the chip may be difficult to describe since it comprises a group of individual components connected together in a particular manner, it is convenient to assume each switch fits into a square of area A , whose side is of length L . Because of the nature of the networks, and the assumption of comparable switch complexity in banyan and crossbar networks, all switches are taken to be of equal size.

The logic associated with the particular switch design can be roughly divided into two parts: One part is associated with data paths (i.e., registers, multiplexers, etc.) while the other is associated with control of these data paths. This control logic for instance would include facilities for local control of the switch position. These parts can be considered to occupy effective square areas A_d and A_c with lengths L_d and L_c respectively. The size of the data portion of the logic is directly related to the number of communication lines associated with each path. This is denoted by w , and is referred to as the path width. This path width, along with other parameters, will directly effect both network bandwidth and chip pin requirements. The latter point will be discussed in a later section. Under many design circumstances it is reasonable to assume that the control portion of the logic is independent of the path width.

A key design parameter, γ , can now be defined as the ratio of A_c to A_d given a path width equal to 1.

$$\gamma = A_c/A_d (w=1) \quad [5]$$

Based on the experience with an SSI crossbar switch design (14,15) a γ on the order of 10 or 20 seems reasonable. This however is a design parameter and may vary widely depending on the control procedures used and resulting logic designed.

The next general assumption relates to the orientation of paths entering and leaving each square switch. It is assumed that one path of width w is associated with each side of the square. While this is clearly a "natural" layout assumption for a crossbar network, some juggling of alternatives indicates it is also reasonable in the banyan case. Note that except for certain switches on the periphery of the network, the others have two entering and two leaving paths.

A second parameter, K , can now be defined as the ratio of the data area A_d to the minimum data area required for a path width equal to w . This minimum area is related to the feature size λ which is determined by the particular technology and fabrication process used. Considering the basic unit of length to be the feature size λ , for

connections laid out in metalization layers, Mead and Conway (17) recommend a minimum line width and minimum distance between adjacent lines of 3 units each. A path containing w parallel communications lines would thus be about $6w$ units wide. Since each switch side is assumed to have w connecting lines, a lower bound on the data portion of the switch area is $(6w)^2$. K may therefore be defined as:

$$K = A_d/(6w)^2 \quad (K \geq 1) \quad [6]$$

K represents the increased area in A_d over the minimum required to support a w wide path.

Substituting the definitions for λ and K in the relationship $A = A_c + A_d$ we now obtain an expression for the switch area in terms of the parameters γ , w , λ and K .

$$A = 36K(\gamma + w^2) \quad [7]$$

The length of a switch side L is therefore:

$$L = 6\sqrt{K(\gamma + w^2)} \quad [8]$$

This concludes the development of a set of general assumptions and definitions related to the switch area. In the following sections these are incorporated in overall topological models for the crossbar and banyan networks.

3.1 A Crossbar Area Model

Figure 4a illustrates a small portion of a larger crossbar network whose overall layout corresponds to that presented in Figure 1. Given a square network with N input and output ports the area required by the network is easily derived as:

$$A_{CB} = [NL + 3(N-1)]^2 \quad [9]$$

3.2 A Banyan Area Model

The banyan VLSI area model is more difficult to derive because, unlike the crossbar, paths between switches cross each other on a single plane, and the area implications of these crossovers are a function of both circuit layout (20), and fabrication technology (e.g., number of connect layers assumed). Determining the "optimal" layout for a given circuit is generally an unsolved problem except for certain simple cases (21). Due to this, the approach taken here is to find the minimum area required by a banyan network, without considering whether a real layout can achieve this minimum. While this biases the result towards the banyan network, indications are that the bias is a minor one and that networks realizations within 30 percent of the minimum are achievable.

With regard to the fabrication technology, the preliminary model assumes two layers of metal interconnect, where one layer provides for all horizontal paths, and the other layer provides for all vertical paths. This can be extended to a three layer model. Further extension to larger numbers of layers are possible, but more difficult to quantify. In addition it is possible to lay out some of the connections in the nonmetalization

layers present. This will change the line width and spacing somewhat but will not change the analysis significantly. For simplicity it is assumed that connections between connect layers (i.e., z direction) take negligible area.

To obtain the banyan network area, consider first the horizontal length of the network. Two possible layouts for a portion of the network are illustrated in Figures 4b and 4c. It is readily seen that the layout in Figure 4b requires less interswitch separation than that in Figure 4c. However, for that to be possible, it is necessary that the vertical spacing occupied by the $2w$ lines between two switches be less than the switch side length. That is, given $2w$ lines each of width 3 units and a separation between lines of 3 units the following condition must hold:

$$12w \leq L = 6\sqrt{K(\gamma+w^2)}$$

Or,

$$2w \leq \sqrt{K(\gamma+w^2)} \quad [10]$$

For example, for a K value of 1.5 and $\gamma=10$, only $w=1$ and $w=2$ satisfy the constraint in [10].

The approach to be taken will be as follows: for w values satisfying [10], expressions will be derived for banyan area based on the layout in Figure 4b. For other values of path width, the layout in Figure 4c will be used. It may be observed that for values of w not satisfying [10], an alternate layout is possible where alternate switches on a row are staggered as in Figure 4d to yield an interswitch separation identical to that in Figure 4b. However, this has the effect of increasing the effective width per row and the results are not significantly different from those obtained following Figure 4c. With a network size N , we obtain the following expressions for the minimum horizontal lengths:

$$\text{For } 2w \leq \sqrt{K(\gamma+w^2)}; L_{BAH} = (N/2)L + 3((N/2)-1)(2w+1) \quad [11a]$$

$$\text{For } 2w > \sqrt{K(\gamma+w^2)}; L_{BAH} = (N/2)L + 3((N/2)-1)(4w+1) \quad [11b]$$

The first term in each equation represents the length contribution due to the switches themselves, while the second term is due to the separation between the switches required by the interconnect paths.

A similar approach can be taken for the vertical length. The vertical separation between switches will be the same for both the banyan layouts. Here, however, the distances between the switches will increase as one moves from one switch row to the next (one level to the next). The reason for this is that the number of path crossovers between succeeding rows increases as one moves from the input ports of the network to the output ports. The result is that the number of adjacent horizontal paths which must be routed increases with each level. Since these paths reside on the same connect layer and must be separated by a minimum distance of 3 units, the

vertical distance between the switches increases correspondingly. Figure 5 illustrates this for a typical second level switch interconnection pattern. The minimum distance between the switch rows can be obtained by taking a vertical cut through the diagram at the point where there is a maximum number of horizontally parallel paths. For the case of $w=1$, $3(2w+1)$ represents the minimum vertical distance, in feature sizes between the switch rows at that level. Notice that this does not depend on the particular routing selected but only on the interconnection topology implicit in the banyan network. With these considerations in mind, the following expression for the total vertical length can be derived.

$$L_{BAV} = L \log_2 N + \sum_{i=2}^{\lceil \log_2(N/2) \rceil} 3n_i + 3 \quad [12]$$

where

$$n_i = 2(2^i - 1)w + 1$$

The first term in 12 represents the length contribution due to the $\log_2 N$ rows of switches. The next two terms represent the sum of the minimum distances achievable between successive rows.

The minimum area required by a banyan network can now be obtained as the product of equations 11 and 12.

$$A_{BA} = L_{BAH} L_{BAV} \quad [13]$$

3.3 Banyan/Crossbar Space Comparisons

To compare the space requirements for the two networks consider the ratio of A_{BA}/A_{CB} . The ratio can be expressed as a function of the design parameters γ, K, w and N . Figure 7 shows this ratio as a function of the network size N , and the path width w . Notice that as the path width increases the area taken by the banyan increases. This is to be expected since, due to crossover considerations in the banyan case, the path width strongly affects row separation. Notice also that the area ratio reaches an asymptotic value as N increases, given by:

$$\lim_{N \rightarrow \infty} \frac{A_{BA}}{A_{CB}} = \frac{3w(\gamma+6w+3)}{(\gamma+3)^2} \quad [14]$$

The interesting point to note here is that this is an entirely different result than that obtained from a traditional cost analysis based on a ratio of number of switches in each network (i.e., $(\log_2 N)/2N$). This traditional approach always shows the banyan being less costly than the crossbar, in addition to the ratio of the costs asymptotically approaching zero for large N . The finite nonzero limit arises from the fact that although the number of switches in the banyan network increases as $N \log_2 N$, the interconnect area increases as N^2 .

4.0 A VLSI Time Delay Model

One approach to determining delay time through a network is to divide this time into two

parts. The first relates to the delay time associated with logic within the switch. The second relates to the delay time encountered in driving the interswitch paths. The sum of these times over an average path constitutes the average delay for a network.

An exact expression for the delay through the switch logic elements requires a detailed knowledge of the switch design. An approximate expression however can be obtained by specifying the number of logic levels (m) a signal passes through in a switch; and by making some specific assumptions about the implementation technology. For our purposes it is reasonable to assume implementation in NMOS with the NOR gate as the typical logic element. If desired, different technology assumptions can be made at this point. This would change the next few equations but would leave the body of the analysis intact. Note also that since the same assumptions are made for both banyan and crossbar networks, it would require major variations in these assumptions to significantly impact the results.

For NMOS NOR gates, the pair gate delay is approximately given (17) as:

$$D_{\text{NOR}} = f(1+q)\tau \quad [15]$$

where f is the gate fanout, τ is the transit time through the active region of an MOS transistor, and q is the ratio of the pullup to pulldown transistor impedances. All gates are assumed to have identical characteristics. A typical design with $q=4$ results in a pair delay of $D_{\text{NOR}} = 5\tau f$. For a switch containing m levels of logic the delay is thus:

$$D_{\text{Sw}} = 5m\tau f/2 \quad [16]$$

This will subsequently be used to determine the total switch delay corresponding to an average network transmission path.

The delays associated with driving any path between switches is a function of both the driving transistor and the capacitance of the driven path. In order to make these delays independent of the path length one design option available is to increase the size of the driving transistor as the path length increases. Another is to increase the number of driver stages per path. These have the effect of reducing path delay at the expense of increased chip area. While such approaches are possible (22), for simplicity, the assumption has been made that all transistors on the chip are of the same size (we omit consideration of the drive transistors at the chip periphery which drive signals off the chip since these would be the same in both the banyan and crossbar cases). The delay associated with interswitch drivers is given (17) by:

$$D_{\text{IS}} = \tau(1+\alpha p) \quad [17]$$

α is the ratio of the interswitch metalization path capacitance per unit area, C_v , to the transistor gate capacitance per unit area, C_g . p is the length of the interswitch connect path to be

driven. Equation 17 will subsequently be used to determine the total path delay corresponding to an average network transmission path.

4.1 A Crossbar Delay Model

As indicated earlier, the average number of switches in a crossbar path is N . Furthermore the crossbar layout discussed in section 3.1 indicated a switch separation equal to three feature sizes (i.e., $p=3$). Substituting in equation 17, and adding in the switch delays expressed in equation 16 results in an overall delay given by:

$$D_{\text{CB}} = 2.5Nm\tau + (N-1)\tau(1+3\alpha) \quad [18]$$

Note the average path in an N by N crossbar contains N switches and $N-1$ interswitch paths.

4.2 A Banyan Delay Model

An equivalent analysis must now be done for the banyan case. Two complicating considerations enter here however, which were not present in the crossbar case. First the path lengths between switches vary both from level to level, and within a given level. Second, the probability of path blocking must be taken into account.

This second point can be handled by directly following the same assumptions which led to equation 1. Thus if P_N is the probability of a message being blocked in a network of size N , and D_{BAS} is the delay associated with a message passing through the network once, then the overall average banyan delay is:

$$D_{\text{BA}} = D_{\text{BAS}}/(1-P_N) \quad [19]$$

D_{BAS} can be expressed in terms of a sum of delays which accrue at each level of the banyan network. Thus, if p_i is the average path length at level i (i.e., between switches on the i^{th} and $(i+1)^{\text{st}}$ rows), D_{BA} can be expressed from equations 16, 17 and 19 as:

$$D_{\text{BA}} = \frac{\tau}{(1-P_N)} \left[2.5m\tau \log_2 N + \sum_{i=1}^{\lceil \log_2(N/2) \rceil} (1+\alpha p_i) \right] \quad [20]$$

The first term represents the contribution due to the $\log_2 N$ rows of switch delays, while the second term represents the contribution due to the $\log_2(N/2)$ interrow path delays. It remains to evaluate p_i . To do this the same connection and minimum distance layout assumptions discussed in the area derivations are utilized.

Two types of paths are present at each level. The first, type 1, is a purely vertical path whose length is $L_{1\text{BAV}}$ given as:

$$L_{1\text{BAV}} = \begin{cases} 3n_i & \text{for } i>1 \\ 3 & \text{for } i=1 \end{cases} \quad [21]$$

where $n_i = 2(2^i-1)w+1$ is the number of horizontal paths in level i . This directly follows the development which led to the vertical distance as

given in equation 12. The second, type 2 path has both horizontal and vertical components. As shown in Figure 6, the vertical component has length of approximately $L_{1BAV} + L/2$. The horizontal component, L_{2BAH} , is given by:

$$L_{2BAH} = (2^{i-1}-1)L + 2^{i-1} \cdot 3(2w+1) + L/2 \quad [22a]$$

for $2w < \sqrt{K(\gamma+w^2)}$

and

$$L_{2BAH} = (2^{i-1}-1)L + 2^{i-1} \cdot 3(4w+1) + L/2 \quad [22b]$$

for $2w > \sqrt{K(\gamma+w^2)}$

The first two terms in each of the expressions above represent the number of switch lengths and switch separations along the type 2 horizontal path. The third term approximates the switch entry point as being in the center of the switch side. Each expression corresponds to an alternative horizontal layout as shown in Figures 4b and 4c.

Given an equal number of type 1 and type 2 paths, and an equal probability of use, p_i is the average of the two type path lengths and can be expressed as:

$$p_1 = (L+6w+9)/2$$

$$p_i = 2^{i-2} (L+30w+3) + 3-6w \quad (i>1) \quad \left. \begin{array}{l} \\ \end{array} \right\} 2w \leq \sqrt{K(\gamma+w^2)} \quad [23a]$$

and

$$p_1 = (L+12w+9)/2$$

$$p_i = 2^{i-2} (L+36w+3) + 3-6w \quad (i>1) \quad \left. \begin{array}{l} \\ \end{array} \right\} 2w > \sqrt{K(\gamma+w^2)} \quad [23b]$$

These equations may be substituted back into [20] to yield the banyan time delay D_{BA} .

4.3 Banyan/Crossbar Delay Comparisons

Corresponding to the area ratio discussed in section 3.3 the delay ratio D_{BA}/D_{CB} can now be formed using equations 18 and 20. Note that this ratio is independent of τ , the transit time. Figure 8 shows this ratio as a function of network size N , and path width w . The curves indicate that as w increases the delay associated with the banyan increases sharply making it worse than the crossbar for values of w greater than four, with the delays being roughly equal for w equal to four. Notice also that as N gets large, for w values greater than 2 the relative delay performance of the crossbar improves. This is due to the fact that the blocking probability now plays a dominant role. Although not obvious from the graph, this will be true for all w values as N gets large enough since $1/(1-P_N)$ increases as N goes to infinity at a rate faster than the other components in the ratio expression decrease. While this is true for both the traditional analysis given in section 2, and this VLSI based analysis, the traditional analysis shows the crossbar delay significantly greater than the banyan delay for the N range given in the graph. This can be seen by the dashed curve in Figure 8.

5.0 Network Space-Time Comparison

Given that banyan and crossbar network area (space) and delay (time) characteristics have been derived, the space time product can be formed using equations 9, 13, 18 and 20.

$$C_{BA} = A_{BA} \cdot D_{BA}; \quad C_{CB} = A_{CB} \cdot D_{CB} \quad [24]$$

From this, the ratio C_{BA}/C_{CB} can be obtained. This ratio is plotted in Figure 9 along with the equivalent ratio as obtained from a traditional analysis [4]. The VLSI analysis shows much less difference in performance than the traditional analysis. In fact a principal conclusion is that for reasonable values of N and current technology parameters the difference in space time products is not significant. As an example, consider the SSI based crossbar switch discussed in references 14 and 15. Each switch requires roughly 2000 gates. Technology projections indicate that VLSI chips containing about 200,000 gates will be available within the next few years. This means that about a 10 by 10 crossbar switch would be possible on a chip. This represents a conservative projection since it is likely that VLSI based switch design would take far fewer gates than the SSI design cited. A path width $w=4$ would require about 80 pins. The area-delay ratio from Figure 9 can be obtained as about .6. Note that the analysis which led to this result made some optimistic assumptions with regard to banyan implementation. In a real design it is likely that the curves would be shifted up somewhat thus indicating even less difference in space-time product performance.

6.0 Summary and Conclusions

This paper has presented a high level VLSI oriented model for comparing banyan and crossbar networks. Comparison has been based on a space-time product performance measure and the assumption of saturated network operation.

The area model, based on general geometric chip layout considerations, was derived from the topological properties of the networks and several layout assumptions. One assumption was that only two connect layers were present, one of them for vertical and the other for horizontal paths. Use of more interconnect layers will reduce the banyan area, however at a significant fabrication cost. The analysis presented may be extended to these cases.

The time model was based on determining the average path of a message through the network. The average number of switches and the average path length were calculated. These results were combined with some simple time delay models based on a NMOS implementation, a typical NOR gate component and factors such as logic fanout and number of levels. Expressions for the average delay through the network were presented.

The area and time models were used to evaluate an overall space-time performance measure for the networks. For reasonable values of N , the difference between the two networks using this measure is relatively small. Given the greater

conceptual and implementation layout simplicity of the crossbar network, it would appear to be the preferable network. Work is now proceeding to see if the analysis techniques presented can be broadened to handle other $N \log N$ networks.

Acknowledgements: The author would like to thank Professor J. R. Cox and Dr. P. N. Turcu for their help in formulating the problem, and K. Padmanabhan for his aid in the subsequent analysis and programming efforts.

References

1. Clos, C., "A Study of Non-Blocking Switching Networks," *The Bell System Technical Journal*, Vol. 32 (March 1953), 406-424.
2. Benes, V. E., Mathematical Theory of Connecting Networks and Telephone Traffic, Academic Press, New York (1965).
3. Waksman, A., "Permutation Networks," *JACM*, Vol. 15, No. 1 (Jan. 1968), 159-163.
4. Batcher, K. E., "Sorting Networks and their Applications," Spring Joint Computer Conference (1968), 307-314.
5. Opferman, D. C. and Tsao-Wu, N. T., "On a Class of Rearrangeable Switching Networks. Part 1: Control Algorithm," *The Bell System Tech. J.*, Vol. 50, No. 5 (May-June, 1971), 1579-1600.
6. Cantor, D. G., "On Non-Blocking Switching Networks," *Networks*, Vol. 1, No. 4 (Winter 1971), 367-377.
7. Lawrie, D. H., "Memory-Processor Connection Networks," Rept. No. UIUCDCS-R-73-557, Dept. of Comp. Sci. Univ. of Illinois at Urbana-Champaign, Urbana, Illinois (1973).
8. Davidson, I. A. and Field, J. A., "Design Criteria for a Switch for a Multiprocessor Computing System," Proc. of the 1975 Sagamore Computer Conference on Parallel Processing, Syracuse University, New York (Aug. 1975), 110-114.
9. Wen, K. Y. and Lawrie, D. H., "Effectiveness of Some Processor/Memory Interconnections," Proc. of 1976 Inter. Conf. on Parallel Processing, Wayne State University, Detroit, Michigan (Aug. 1976), 283-292.
10. Pearce, R. C. and Majithia, J. C., "Upper Bounds on the Performance of Some Processor-Memory Interconnections," Proc. of the 1976 Inter. Conf. on Parallel Processing, Wayne State University, Detroit, Michigan (Aug. 1976), 303.
11. Siegel, H. J., "Analysis Techniques for SIMD Machine Interconnection Networks and Effects of Processor Address Masks," *IEEE T-EC*, Vol. C-26, No 2 (Feb. 1977), 153-161.
12. Patel, J. H., "Processor Memory Interconnections for Multiprocessors," *The 6th Annual Symposium on Computer Architecture* (April 1979), 168-177.
13. Pippenger, N., "On Crossbar Switching Networks," *IEEE T-COM*, Vol. COM-23, No. 6 (June 1975), 646-659.
14. Kahn, S. H. "Design of a Modular Crossbar Network for a Multiprocessor," M.S. Thesis, Dept. of Elec. Engr., Washington Univ., St. Louis (1978).
15. Franklin, M. A., Kahn, S. A. and Stucki, M. J., "Design Issues in the Development of a Modular Multiprocessor Communications Network, The 6th Ann. Symp. on Comp. Arch. (April 1979), 182-187.
16. Goke, L. R. and Lipovski, G. J., "Banyan Networks for Partitioning Multiprocessor Systems," *The First Ann. Symp. on Comp. Arch., University of Florida, Gainesville, Florida* (1973), 21-28.
17. Mead, C. and Conway, L., Introduction to VLSI Systems, Addison-Wesley (1979).
18. Loneragan, W. and King, P., "Design of the B5000 System," In Computer Structures: Readings and Examples, Bell, C. G. and Newell, A., McGraw-Hill (1971).
19. Wulf, W. A. and Bell, C. G., "C.mmp--A Multi-Mini-Processor," *Proc. AFIPS Conf.* (Fall 1972).
20. Turcu, P. N., "Modelling Analysis of Switch Interconnection Networks," Ph.D. Thesis, Dept. of Comp. Sci., Washington Univ., St. Louis (1979).
21. Cutler, M. and Shiloach, Y., "Permutation Layout," *Networks*, Vol. 8 (1978), 253-278.
22. Mead, C. A. and Rem, M., "Cost and Performance of VLSI Computing Structures," *IEEE Proc. of Solid-State Ckts.*, Vol. SC-14, No. 2 (April 1979).




FIGURE 1A: 8 * 8 CROSSBAR SYSTEM




FIGURE 2A: 8 * 8 BANYAN SYSTEM




FIGURE 1B: CROSSBAR SWITCH POSITIONS




FIGURE 2B: BANYAN SWITCH POSITIONS




FIGURE 3: MESSAGE BLOCKING PROBABILITY (P_M) VERSUS NETWORK SIZE ($\log_2 N$) (FOR BANYAN NETWORKS WITH UNIQUE AND UNIFORMLY DISTRIBUTED DESTINATION ADDRESSES)




FIGURE 4 A: DETAIL OF CROSSBAR LAYOUT GEOMETRY
 B: MINIMUM HORIZONTAL LAYOUT GEOMETRY FOR BANYAN [$w=1, 2\lambda \sqrt{k(\gamma+w^2)}$]
 C: MINIMUM HORIZONTAL LAYOUT GEOMETRY FOR BANYAN [$w=1, 2\lambda \sqrt{k(\gamma+w^2)}$]
 D: ALTERNATE LAYOUT TO FIGURE 4C (STAGGERED ARRANGEMENT)




FIGURE 5: LEVEL TWO MINIMUM VERTICAL LAYOUT GEOMETRY FOR BANYAN ($w=1$) (ONLY PATHS WITH HORIZONTAL COMPONENT SHOWN)




FIGURE 6: TYPICAL LEVEL 3, TYPE 1 AND TYPE 2 SWITCH CONNECTION PATHS




FIGURE 7: RATIO OF BANYAN TO CROSSBAR AREA REQUIREMENTS (DASHED CURVE IS TRADITIONAL ANALYSIS)




FIGURE 8: RATIO OF BANYAN TO CROSSBAR TIME DELAYS (DASHED CURVE IS FROM TRADITIONAL ANALYSIS)




FIGURE 9: RATIO OF BANYAN TO CROSSBAR SPACE TIME PRODUCTS (DASHED CURVE IS FROM TRADITIONAL ANALYSIS)