

COMPACT LAYOUTS OF BANYAN/FFT NETWORKS\*

David S. Wise

Computer Science Department  
Indiana University  
Bloomington, IN 47405

TECHNICAL REPORT NO. 115  
COMPACT LAYOUTS OF BANYAN/FFT NETWORKS

DAVID S. WISE

AUGUST, 1981

\*To be presented at CMU Conference on VLSI Systems and Computations (Proceedings to be published by Computer Science Press), October, 1981.

This research was supported in part by the National Science Foundation under grant number MCS 77-22325.

## COMPACT LAYOUTS OF BANYAN/FFT NETWORKS\*

David S. Wise

Computer Science Department  
Indiana University  
Bloomington, IN 47405

ABSTRACT A two-layer pattern is presented for the crossover pattern that appears as the FFT signal flow graph and in many switching networks like the banyan, delta, or omega nets. It is characterized by constant path length and a regular pattern, providing uniform propagation delay and capacitance, and ameliorating design problems for VLSI implementation. These are important issues since path length grows linearly with the number of inputs to such networks, even though switching delay seems to grow only logarithmically.

Independent of that pattern, an arrangement for stacking planes of such planar crossover patterns in three dimensions is described. Whereas a planar crossover pattern of  $\theta(m)$  inputs and outputs has at best  $O(m)$  path length, the stacked pattern allows  $O(\sqrt{m})$  path length. The scheme provides for stacking  $2k$  planar nets (perhaps VLSI chips), each with  $k$  inputs/outputs into a network of  $k^2$  inputs/outputs. Using this pattern, all such paths would have length (propagation delays) =  $O(k)$ .

Key Words and Phrases: VLSI, switching networks, delta networks, omega networks, multiprocessing.

CR Categories: 6.1, 4.32, 3.81, 3.83.

### I. INTRODUCTION

This paper offers two results that can both be described as pictures. They are Figures 1 and 3. The perceptive reader may stop here, since the remainder of this paper only describes them.

---

\*This research was supported in part by the National Science Foundation under grant number MCS 77-22325.



## II. NOTATION

The logarithm-base-2 of  $n$  is written  $\lg x$ . For real functions  $f$  and  $g$ , we write  $f(x) = O(g(x))$  if there is a constant  $k$  and some value  $x_0$  such that  $f(x) \leq k \cdot g(x)$  for all  $x > x_0$ . This notation expresses a proportional, asymptotic upper bound of  $g$  for  $f$ . If  $f(x) = O(g(x))$  and  $g(x) = O(f(x))$  then we write  $f(x) = \Theta(g(x))$ . It is only necessary to express  $g$  as a single term (e.g.  $\lg x$ ,  $2^x$ ,  $x^2$ ) in such contexts.

The abbreviations FFT and VLSI refer to the Fast Fourier Transform and Very Large Scale Integration circuit technology, respectively.

## III. AN APPLICATION

This work is immediately motivated by the need for a switching network between processors and memory in a multi-processor system of 100 or 1000 processors [4]. In order to increase bandwidth to memory, reducing contention among the processors, a banked memory is envisioned. Its access is through a fast, parallel switching network.

A suitable model for such a network is a banyan network [7] whose elementary functional unit is a  $2 \times 2$  crossbar switch. It may be perceived as a router [2], a store-and-forward unit with two input lines and two output lines. Figure 1 might be interpreted as such a network from  $2n = 16$  processor to  $2n$  memories.

Memory fetch and store instructions are transmitted as packets through the network; duplicate networks pass information in the reverse direction. (Say, honoring a fetch instruction or allocating free nodes from a heap [5].) Each packet initially contains a binary (destination) memory-address followed by a message. Upon arrival at each router, its high-order address list determines along which path it is to be forwarded. The entire message is shifted left one bit, displacing that address bit for transmission to the next stage. In Figure 1, a zero bit would send the modified packet from a router on a leftward (northwest) line; a one bit would send the remainder of the packet to the right (northeast). It is possible to insert into the vacated low-order bit a value identifying the input line by which a packet entered each router. After a packet has passed through the network to its destination, its destination address will have been shifted out. In its stead (at the end of the packet) would be the address of its source processor when the vacant bits are so used.

This describes a network which uses crossover pattern to allow as many as  $n$  messages to arrive at some of  $2n$  different destinations simultaneously, each over a path of at most  $\lg n$  routers. As many as  $n \cdot \lg n$  messages might be already in the switch, pipelined behind the arriving wave.



On the other hand, contention is still possible (for instance, all messages going to one destination) since each router can only handle one message at a time.

#### IV. PLANAR LAYOUT WITH CONSTANT PATH-LENGTH.

Figure 1 presents an alternative to the wiring diagram which is the essential component of many switching networks, such as the banyan network, the delta network, the omega network, and the Benes network (as either a circuit switch or a packet switch) [6,9]. It also appears as the broadcast pattern for intermediate results in the FFT [1].

The functional boxes, pictured as shaded boxes in the figure, are of little importance here. Their definitions vary with the various applications of this crossover network. The only assumption we make is that they are all of uniform design, perhaps mirror images of one another. Their size affects the aggregate size of the network.

Most analyses of such networks emphasize these functional boxes as the critical cost of such networks. That is, any path from the  $n$  input boxes (at the bottom) through the  $n$  output boxes (at the top) of such a graph passes through exactly  $\lg n$  such boxes. If such a network is implemented in VLSI technology, however,  $n$  can become locally large and the apparently low  $O(\log n)$  delay may be unattainable when timing considerations due to line-delay are included [3]. While Figure 1 has the property that all lines at any stage are of equal length, it still exhibits this line delay-problem. Drawn with uniform inter-wire spacing, it shows that the spacing between the stages, between functional boxes along any path, grows exponentially (top-to-bottom). Thus, although there are only  $\lg n$  stages, every path length is  $\Theta(n)$ .

For any planar graph with  $n$  inputs/outputs, the longest path length is  $kn$ , for some constant  $k$ . This is easily proved, since the upper-left output must have a path from the lower-right input, and the network has width  $n$ . Figure 1 exhibits the desirable property that all paths are of length  $kn$ , so that special considerations due to line delay, attenuation, and capacitance are lessened. Functional boxes may be defined more uniformly, a desideratum in VLSI design.

Also important for planar fabrication technology is the fact that the crossover network is exactly two layers thick. While it exhibits more wire crossings than the common planar patterns (e.g. the cover of [1]), the others require three or more layers. Here one layer is composed of all diagonal wires running "northeast" (lower-left to upper-right) and the other layer is composed of "northwest" wires (lower-right to upper-left). The northeast wires might be in the metal layer of a VLSI chip and the northwest wires in the diffusion or (better) a second metal






Figure 1. Crossover Net for  $n=8$

layer. Alternatively the northeast wires might be on one side of a printed circuit board and the northwest wires on the other. The D-shaped blobs in the figure indicate contact cuts through the insulator connecting the layers.

The size of the contact-cut is one constraint on the height of the network. According to Mead and Conway, [7] their diameter, including insulating spacing, is  $7\lambda$  in VLSI, larger than any insulated wire. Let that wire diameter be  $w$ , and let us measure the width of one functional box, including insulation as a multiple of  $w = qw$ .\* As Figure 2 illustrates, the angle of either diagonal is then  $\arcsin 1/q$ . This is another constraint on the height of the network.

In fact, the height of the  $n$  input/output crossover net of Figure 1 can then be calculated accurately as

$$\sum_{0 \leq i < \lg n} 2^i \cdot qw \cdot \tan \theta$$

$$= w(n-1)q/\sqrt{q^2-1}.$$

To this must be added the height of  $(\lg n + 1)$  functional boxes.

The path length along any path is the height of the network times  $q = \operatorname{cosec} \theta$  plus the path within  $(\lg n + 1)$  functional boxes. This is not quite accurate, since additional width of the crossover net is necessary to prevent contact cuts from overlapping. Over all, the width must be increased by  $wn/4$  to allow for the  $n/4$  pairs of cuts across the second level of Figure 1; the top level can be wired without contact cuts.

It is then possible to derive the exact area of the crossover network. The width is  $nw(q+0.25)$  and the area of the wiring is  $w^2q(q+0.25)(q^2-1)^{-0.5}(n^2-n) = \theta(n^2)$ . To this must be added  $w(q+0.25)n(\lg n+1)$  times the height of one functional box. This can be compared with the asymptotic area of a shuffle-exchange network [6] of  $\theta((n/\log n)^2)$ ; that network performs unpipelined FFT or packet switching (not circuit switching) with  $n$  functional boxes in  $(\lg n+1)$  iterations, accounting for one factor of  $\log n$  in its compression of area.

## V. RACKING NETWORKS IN CUBIC SPACE.

Regardless of the technology or design of a planar crossover network of the sort presented above, extending it to large  $n$  -- say  $10^4$  or  $10^6$  -- requires a large plane. As we have seen, and as Franklin [3] points out, both width

---

\*This does not allow horizontal spacing for immediately adjacent contact cuts -- the facing D's in Figure 1 -- but we are measuring vertical distance.






Figure 2.  $\Theta = \text{Arcsin } 1/q.$

and height grows linearly with  $n$ . Obviously some dicing of larger nets will be required as  $n$  increases.

Figure 3 illustrates an elegant decomposition from the plane into cubic space. It is elegant because of its decomposition into planes, and because of a net reduction in the height of the network without increasing its volume. Built of  $4n$  planar nets, each  $n$  functional boxes wide, it is the equivalent of a network with  $2n^2$  functional boxes on either edge. The planar nets are arranged into two racks of  $2n$  planes each; the racks are stacked one atop the other but with orthogonal racking arrangements.


The interconnection between the racks is very simple: there are no wire crossings there; all wires run nearly vertically. Therefore, there is no required distance between the racks, except a constant spacing that depends on the technology used to rack the planes. Thus, the height of such a network is only about twice the height of one planar component, linear in  $n$  as discussed above. The height of a network that is  $2n^2$  functional boxes "wide" is only  $O(n)$  tall. More importantly the path length is also only  $O(n)$  instead of the  $O(n^2)$  of equivalent planar networks.

The two outputs from one of the final functional boxes on a lower plane are routed to corresponding initial boxes on each of two adjacent upper planes; the dual perspective is that inputs to one functional box on an upper plane come from corresponding output boxes on each of two adjacent lower planes.

That such a pattern is the equivalent to extensions of Figure 1 (that have  $2^{2k+1}$  boxes across the bottom and top) can be seen from the following argument. Split such a network by temporarily removing the middle crossover web. The top half is  $2^{k+1}$  disjoint planar networks of  $2^k$  boxes in width. Although the wires cross, the bottom half may be untangled into exactly the same pattern. (In Figure 1, for instance, the first and last boxes in the bottom row are connected to the middle two boxes in the next row above, just as the left two boxes in the top row are connected with those immediately under them.) If we visualize the  $2^{k+1}$  networks in the top half as stacked but with every other "card" in the stack reversed in mirror image, and if the  $2^{k+1}$  networks in the bottom half are untangled and stacked, then the missing middle web is the analog of the vertical connections described above.

Using the switching example of Section III, this racking can be perceived a different way. The purpose of the entire network is to switch a packet through to a destination which is at an  $(x,y)$  co-ordinate in the top plane. Its address is a bit concatenation of  $x$  and  $y$ . The lower rack uses the high order ( $x$ ) bits of the address to route the packet to a correct position on a planar circuit oriented along the  $x$ -axis. Then it passes vertically onto an identical board oriented along the  $y$ -axis, which uses





**Figure 3.** Racking of Planar  $n$  Networks in Cubic Space.

the low order bits of the address to forward it to its correct destination.

This example points out an application of the racking in which all  $2^n$  planar circuits, including functional boxes, are identical, making it quite attractive for planar fabrication techniques, such as those used in current VLSI or printed-circuit-board technology. Alternatively it shows how existing planar circuits for FFT or switching can be used in racks to square their effectiveness.

## VI. CONCLUSION.

This paper shows how to lay out a crossover network either on a plane in just two layers and with uniform wire length or in cubic space in order to minimize path length. Since the latter pattern (Figure 3) is fabricated from planar nets, the former results apply for it as well.

Although the two solutions have a common motivation, they may be used independently of one another, either in a case where a planar layout is required (Figure 1) or where a planar net is readily available but an extension is needed.

A feature of both solutions is that no connections to function boxes are made horizontally. In Figure 1, they are accessible from the sides for busses across each stage. Common signals -- like power busses -- can be passed across the boxes without interfering with the crossover net under consideration. In Figure 3 the racked planes run completely across the circuit, leaving the edges available for maintenance lines, and gaps between the planes (in the case of printed circuit technology) for cooling.

It is possible to build up an arrangement equivalent to Figure 3 from smaller planar nets (say identical to Figure 1) using the recurrence of Figure 3 repeatedly. The resulting arrangement would not have the features of simple stacking just mentioned, but it might be necessary if fabrication techniques could not provide sufficiently large planar circuits, and might be useful when technologies must be mixed as suggested below.

Fabrication technology is a constraint to importance of the asymptotic analysis of these circuits. In many applications, like Section III above, each line in the figures represents a parallel communication bus of from 64 to 256 wires. Large planar circuits are not now available from VLSI techniques when the number of such input/output lines is high. The inter-rack connection of Figure 3, though simple, is not now possible in the same scale as the VLSI planar network. The "short" vertical wires between the racks would be very, very long today. Thus the importance of Figure 3 for planar circuits on a chip is now conceptual rather than practical; path length is probably not reduced under current fabrication techniques, but the design is still a useful modulization of Figure 1. Thus,



we might well use Figure 1 to lay out a small planar VLSI chip (say  $n=8$ ), then use Figure 3 to build a printed circuit board with 32 such chips ( $n=128$ ), and finally rack such boards using Figure 3 to build a very large network ( $n=32768$ ). The path lengths would not be significantly reduced by using these designs on one board, but the regular pattern on a chip, and within the rack of boards (instead of a backplane) would make the aggregate speed of the circuit faster and easier to clock than conventional planar layouts of that size.

ACKNOWLEDGEMENT: I thank John O'Donnell for introducing me to the depth of VLSI technology, George Epstein for suggesting [1], C. E. Leiserson for pointing out [6], and Nancy Garrett, able typist.

## REFERENCES

1. N. Ahmed and K.R. Rao. Orthogonal Transforms for Digital Signal Processing. Berlin, Springer (1975).
2. J.B. Dennis, G.A. Boughton, and C.K.C. Leung. Building blocks for data flow prototypes. Proc. 7th Symp. Computer Architecture (IEEE order no. 80CH 1494-4C), ACM SIGARCH Newsletter 8, 3 (May, 1980), 1-8.
3. M.A. Franklin. VLSI performance comparison of banyan and crossbar communication networks. IEEE Trans. Computers C-30, 4 (April, 1981), 283-291.
4. D.P. Friedman and D.S. Wise. Aspects of applicative programming for parallel processing. IEEE Trans. Comput. C-27, 4 (April, 1978), 289-296.
5. S.D. Johnson. Connection networks for output-driven list multi-processing. Technical Report 114, Computer Science Department, Indiana University (October, 1981).
6. D. Kleitman, F.T. Leighton, M. Lepley, and G.L. Miller. New layouts for the shuffle-exchange graph. Proc. 13th ACM Symp. on Theory of Computing, [ACM Order No. 508810] (1981), 278-292.
7. G.M. Masson, G.C. Gingher, and S. Nakamura. A sampler of circuit switching networks. Computer 12, 6 (June, 1979), 32-48.
8. C. Mead and L. Conway. Introduction to VLSI Systems, Reading, MA, Addison-Wesley (1980).
9. H. Siegel. Interconnection networks for SIMD machines. Computer 12, 6 (June, 1979), 57-65.