

CHAPTER 1

FOUNDATIONS

In this chapter, we define the shuffle-exchange graph and the Thompson grid model of a chip. We also review the previous work on layouts for the shuffle-exchange graph. In Section 1.3, we describe Thompson's straightforward $O(N^2/\log^{1/2}N)$ -area layout for the N -node shuffle-exchange graph, and in Section 1.4, we describe Hoey and Leiserson's complex plane diagram. The complex plane diagram is useful for finding good layouts of the shuffle-exchange graph. For example, Hoey and Leiserson used the diagram to find an $O(N^2/\log N)$ -area layout for the N -node shuffle-exchange graph in [34]. In Chapter 2, we will use the diagram to find a variety of layouts for the N -node shuffle-exchange graph including one that requires only $O(N^2/\log^{3/2}N)$ area. The complex plane diagram will also be used in Chapter 4 as an aid in the construction of good practical layouts for small shuffle-exchange graphs.

1.1 The Shuffle-Exchange Graph

The *shuffle-exchange graph* comes in various sizes. In particular, there is an N -node shuffle-exchange graph for every N which is a power of two. Each node of the $(N=2^k)$ -node shuffle-exchange graph is associated with a unique k -bit binary string $a_{k-1} \cdots a_0$. Two nodes w and w' are linked via a *shuffle edge* if w' is a left or right cyclic shift of w (i.e., if $w = a_{k-1} \cdots a_0$ and $w' = a_{k-2} \cdots a_0 a_{k-1}$ or

$w' = a_0 a_{k-1} \dots a_1$, respectively). Two nodes w and w' are linked via an *exchange edge* if w and w' differ only in the last bit (i.e., if $w = a_{k-1} \dots a_1 0$ and $w' = a_{k-1} \dots a_1 1$ or vice-versa). As an example, we have drawn the 8-node shuffle-exchange graph in Figure 1-1. Note that the shuffle edges are drawn with solid lines while the exchange edges are drawn with dashed lines. We shall follow this convention throughout the book.




Figure 1-1: The 8-node shuffle-exchange graph.

By replacing the nodes and edges of the shuffle-exchange graph by processors and wires (respectively), the shuffle-exchange graph can be transformed into a very powerful parallel computer (which we call the *shuffle-exchange computer*). The computational power of the shuffle-exchange computer is partly derived from the fact that every pair of nodes in an N -node shuffle-exchange graph is linked by a path containing at most $2\log N$ edges and thus the communication time between any pair of processors is short.

More importantly, however, the shuffle-exchange computer is capable of performing a perfect shuffle on a set of data in a single parallel operation. For example, consider a deck of 8 cards distributed among the 8 processors of the 8-node shuffle-exchange graph so that processor 000 initially has card 0, processor 001 initially has card 1, processor 010 initially has card 2, and so forth. Next,

consider a (parallel) operation of the shuffle-exchange computer in which each processor $a_2a_1a_0$ sends its card across a shuffle edge to the neighboring processor $a_1a_0a_2$. It is easily verified that, after completion of the operation, processor 000 contains card 0 (the top card in the shuffled deck), processor 001 contains card 4 (the second card in the shuffled deck), and so forth.

The power of card shuffling and its mathematical abstractions are well known to magicians and mathematicians [23] as well as to computer scientists [87, 90]. For a good survey of the computational power of the shuffle-exchange graph, we recommend Schwartz' paper on ultracomputers [87]. In addition, Stone's paper [90] contains a nice description of some important parallel algorithms based on the shuffle-exchange graph.

1.2 The Thompson Grid Model

Among the many mathematical models that have been proposed for VLSI computation, the most widely accepted is due to Thompson and is known as the *Thompson grid model* [92, 93]. The grid model of a VLSI chip is quite simple. The chip is presumed to consist of a grid of vertical and horizontal *tracks* which are spaced apart by unit intervals. Processors are viewed as points and are located only at the intersection of grid tracks. Two layers of interconnect are used to route the wires. Vertical wires are routed in the top layer of interconnect and horizontal wires are routed in the bottom layer. Hence wires may cross but cannot overlap for any distance. Nor can wires overlap processors to which they are not adjacent. *Contact cuts*, which connect segments of the same wire which are in different layers, are also only located at the intersection of grid tracks. (The routing of wires in this fashion is also known as *layer per direction routing* and *Manhattan routing*.)

As an example, we have included a grid layout for the 8-node shuffle-exchange graph in Figure 1-2. As before, the shuffle edges are drawn with solid lines while

the exchange edges are drawn with dashed lines. Although we have not included them in the figure, contact cuts would be placed at points where a wire changes direction. Also notice that we have omitted the self-loops in Figure 1-2 since they are electrically redundant. In general, the processors need not all be placed on a single horizontal line (as they are in this example).




Figure 1-2: *A grid model layout of the 8-node shuffle-exchange graph.*

Practical considerations dictate that the area of a VLSI layout be as small as possible. The *area of a layout* in the grid model is defined to be the product of the number of horizontal tracks and the number of vertical tracks which contain a processor or wire segment of the layout. For example, the layout in Figure 2 has area 48. As can be easily observed, this is far from optimal.

Other measures of interest are the wire area, crossing number, maximum edge length and maximum edge crossing of a layout. The *wire area* of a layout is defined to be the sum of the lengths of the wires in the layout. The *crossing number* is the number of points in the layout where two wires cross. The *maximum edge length* is the length of the longest wire in the layout. The *maximum edge crossing* is the maximum number of points at which a single wire crosses other wires. For example, the layout in Figure 1-2 has wire area 50, crossing number 8,

maximum edge length 9 and maximum edge crossing 4. The reader can verify that none of these values is optimal by skipping ahead to Figure 4-3.

1.3 An $O(N^2/\log^{1/2}N)$ -Area Layout

Thompson was the first to investigate VLSI layouts for the shuffle-exchange graph. In his thesis [93], he showed that *any* layout for the N -node shuffle-exchange graph requires at least $\Omega(N^2/\log^2N)$ wire area. (We reprove this fact using crossing number arguments in Part II of the thesis.) In addition, he described a layout requiring only $O(N^2/\log^{1/2}N)$ area. In what follows, we present Thompson's layout and give a simple proof that it does, in fact, require just $O(N^2/\log^{1/2}N)$ area.

Given any k -bit string w , define the *size* of w to be the number of 1-bits it contains. For example, the size of 10110 is 3. Thompson's idea was to lay out the $N=2^k$ nodes of the shuffle-exchange graph on a straight line in order of nondecreasing size. It is easily seen that shuffle edges link nodes which have the same size and that exchange edges link nodes which have sizes differing by one. Thus the edges of such a layout are relatively short. In fact, nodes connected by shuffle edges can be placed in groups, so that only 2 horizontal tracks are used for all the shuffle connections. The remaining horizontal tracks are occupied by exchange edges.

The exchange edges are inserted from left to right so that each exchange edge occupies two vertical tracks and a portion of the lowest horizontal track which is empty at the time of its insertion. (For example, Figure 1-2 displays a layout for the 8-node shuffle-exchange designed in this way.) This well-known strategy for inserting exchange edges guarantees that the number of horizontal tracks used will be minimal, and equal to the maximum number of edges which must (at some fixed point) overlap one another. Since exchange edges link nodes which differ in

size by one, it is easily seen that the maximum overlap is at most $O(\max_{0 \leq s \leq k} B_s)$ where B_s is the number of nodes of size s .

It is easy to show that $B_s = C(k, s)$ for each s , where

$$C(k, s) = k! / [s!(k-s)!]$$

is the well-known function for binomial coefficients. It is also well-known that $C(k, s)$ achieves its maximum value at $s = k/2$ for any k . Using standard asymptotic analysis, it is easily shown that $C(k, k/2) \sim (2/\pi)^{1/2} (2^k / k^{1/2})$ for large k . (For a good review of such techniques, see Bender and Orszag's book [6].) Thus Thompson's layout requires only $O(N/\log^{1/2} N)$ horizontal tracks. Since only 1 or 2 vertical tracks are needed to embed the vertical portions of the edges incident to any given node, we can conclude that Thompson's layout has area $O(N^2/\log^{1/2} N)$.

1.4 The Complex Plane Diagram

In [34], Hoey and Leiserson observed that there is a very natural embedding of the shuffle-exchange graph in the complex plane. In what follows, we describe this embedding (henceforth referred to as the *complex plane diagram*) and point out some of its more important properties. In addition, we give a brief description of the method used by Hoey and Leiserson to transform the diagram into an $O(N^2/\log N)$ -area layout for the N -node shuffle-exchange graph.

1.4.1 Definition

Let $\delta_k = e^{2\pi i/k}$ denote the k th primitive root of unity. Given any k -bit binary string $w = a_{k-1} \dots a_0$, let $p(w)$ be the map which sends w to the point

$$p(w) = a_{k-1}\delta_k^{k-1} + \dots + a_1\delta_k + a_0$$

in the complex plane. As each node of the $(N = 2^k)$ -node shuffle-exchange graph corresponds to a k -bit binary string, it is possible to use the map to embed the

shuffle-exchange graph in the complex plane. For example, we have done this for the 32-node shuffle-exchange graph (whence $k=5$) in Figure 1-3. As usual, we have drawn the shuffle edges with solid lines and the exchange edges with dashed lines. For simplicity, each node is labeled with its value instead of its 5-bit binary string. (By the *value* of a node, we mean the numerical value of the associated k -bit binary string. For example, the value of *01101* is 13.)




Figure 1-3: *The complex plane diagram for the 32-node shuffle-exchange graph. (Taken from [34].)*

1.4.2 Properties

Examination of Figure 1-3 indicates that the complex plane diagram has some very interesting properties. First, it is apparent that the shuffle edges occur in cycles (which we call *necklaces*) which are symmetrically placed about the origin.

This phenomenon is easily explained by the following identity:

$$\begin{aligned}
 \delta_k p(a_{k-1} \cdots a_0) &= a_{k-1} \delta_k^k + a_{k-2} \delta_k^{k-1} + \cdots + a_1 \delta_k^2 + a_0 \delta_k \\
 &= a_{k-2} \delta_k^{k-1} + \cdots + a_0 \delta_k + a_{k-1} \\
 &= p(a_{k-2} \cdots a_0 a_{k-1}).
 \end{aligned}$$

Thus traversal of a shuffle edge corresponds to a $2\pi/k$ rotation in the complex plane.

Except for degenerate cases, the preceding identity also indicates that each necklace contains k nodes, each a cyclic shift of the other. Such necklaces are called *full necklaces*. *Degenerate necklaces* contain fewer than k nodes and, because they must have some symmetry, are mapped entirely to the origin of the complex plane diagram. For example, $\{00000\}$ and $\{0101, 1010\}$ are degenerate necklaces while both $\{101, 011, 110\}$ and $\{11100, 11001, 10011, 00111, 01110\}$ are full. As we note in the following lemma, the number of degenerate necklaces is quite small compared to the number of full necklaces.

Lemma 1-1: *There are $O(N^{1/2})$ degenerate necklaces and $N/\log N - O(N^{1/2}/\log N)$ full necklaces in the N -node shuffle-exchange graph.*

Proof: A node w is in a degenerate necklace if its binary representation has a nontrivial symmetry with respect to cyclic shifts. Without loss of generality, such a string of bits must consist of a block of k/p bits which is repeated p times where p is some prime divisor of k . As there are $2^{k/p}$ binary strings of length k/p , this means that the number of nodes in degenerate necklaces is at most

$$\sum_{p \geq 2}^{pk} 2^{k/p} \leq O(N^{1/2}).$$

The remaining $N - O(N^{1/2})$ nodes are in full necklaces. As each full necklace contains $\log N$ nodes, there are $N/\log N - O(N^{1/2}/\log N)$ full necklaces. $\square$

It will often be convenient to refer to a necklace by one of its nodes. In particular, we will use the notation $\langle w \rangle$ to indicate the *necklace generated by w* . This is simply the collection of cyclic shifts of w . For example, the necklace generated by 101 is $\langle 101 \rangle = \{101, 011, 110\}$.

Exchange edges are also embedded in a very regular fashion in the complex plane diagram. In fact, each exchange edge is embedded as a horizontal line segment of unit length. This phenomenon is explained by the identity

$$\begin{aligned} p(a_{k-1} \dots a_1 0) + 1 &= a_{k-1} \delta_k^{k-1} + \dots + a_1 \delta_k + 1 \\ &= p(a_{k-1} \dots a_1 1). \end{aligned}$$

In some cases, several exchange edges are contained in the same horizontal line of the diagram. Such lines are called *levels*. For example, there are 9 levels in the diagram of the 32-node shuffle-exchange graph shown in Figure 1-3. We will use the properties of levels in Chapter 2 to find an $O(N^2/\log^{3/2} N)$ -area layout for the N -node shuffle-exchange graph. They will also be used in Chapter 4 to find good practical layouts for small shuffle-exchange graphs.

1.4.3 An $O(N^2/\log N)$ -Area Layout

In [34], Hoey and Leiserson show how to use the complex plane diagram to construct an $O(N^2/\log N)$ -area layout for the N -node shuffle-exchange graph. Their method is quite complicated, however, and we have chosen not to include it here. The basic idea is to use the structural properties of the complex plane diagram to find an $O(N/\log^{1/2} N)$ -separator for the N -node shuffle-exchange graph whenever N is of the form 2^{2^r} for some $r \geq 0$. The separator can then be used to construct an $O(N^2/\log N)$ -area layout by using Leiserson and Valiant's general layout technique for graphs with known separators. (Separators and their application to layouts are discussed in Part II.)

Shortly after writing [34], Hoyer and Leiserson found a far simpler $O(N^2/\log N)$ -area layout for the N -node shuffle exchange graph which was, in addition, valid for all N . By the that time, however, we (as well as several others) had also observed that the complex plane diagram could be used to find a simple layout for the shuffle-exchange graph. This layout is described in Chapter 2.