

LABORATORY FOR
COMPUTER SCIENCE



MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

MIT/LCS/TM-221

LAYOUTS FOR THE SHUFFLE-EXCHANGE GRAPH
BASED ON THE COMPLEX PLANE DIAGRAM

Frank Thomson Leighton
Margaret Lepley
Gary L. Miller

June 1982

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

LAYOUTS FOR THE SHUFFLE-EXCHANGE GRAPH
BASED ON THE COMPLEX PLANE DIAGRAM

Frank Thomson Leighton
Margaret Lepley
Gary L. Miller

Mathematics Department
and
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Abstract: The shuffle-exchange graph is one of the best structures known for parallel computation. Among other things, a shuffle-exchange computer can be used to compute discrete Fourier transforms, multiply matrices, evaluate polynomials, perform permutations and sort lists. The algorithms needed for these operations are extremely simple and many require no more than logarithmic time and constant space per processor. In this paper, we analyze the algebraic structure of the shuffle-exchange graph in order to find area-efficient embeddings of the graph in a two-dimensional grid. The results are applicable to the design of Very Large Scale Integration (VLSI) circuit layouts for a shuffle-exchange computer.

Key words: area-efficient chip layouts, complex plane diagram, graph embedding, necklace, shuffle-exchange graph, Thompson grid model, Very Large Scale Integration (VLSI)

This research was supported in part by National Science Foundation Grant MCS 80-07756 and Defense Advanced Research Projects Agency Grant N00014-80-C-0622.

LAYOUTS FOR THE SHUFFLE-EXCHANGE GRAPH BASED ON THE COMPLEX PLANE DIAGRAM

Frank Thomson Leighton
Margaret Leiby
Gary L. Miller

Mathematics Department
and
Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Abstract: The shuffle-exchange graph is one of the best structures known for parallel computation. Among other things a shuffle-exchange computer can be used to compute discrete Fourier transforms, multiply real and complex polynomials, perform permutations and sort lists. The algorithms needed for these operations are extremely simple and many require no more than logarithmic time and constant space per processor. In this paper we analyze the algebraic structure of the shuffle-exchange graph in order to find more efficient embeddings of the graph in a two-dimensional grid. The results are applicable to the design of Very Large Scale Integration (VLSI) circuit layouts for a shuffle-exchange computer.

Key words: more efficient chip layouts, complex plane diagram, graph embedding, necklace, shuffle-exchange graph, Thompson grid method, Very Large Scale Integration (VLSI)

This research was supported in part by National Science Foundation Grant MCS 80-07356 and Defense Advanced Research Projects Agency Grant N00014-80-C-0613.

1. Introduction

The shuffle-exchange graph has long been recognized as one of the best structures known for parallel computation. Among its many applications, a shuffle-exchange computer can be used to compute discrete Fourier transforms, multiply matrices, evaluate polynomials, perform permutations and sort lists [S71, P80, S80]. The algorithms needed for these operations are extremely simple and many require no more than logarithmic time and constant space per processor.

Recent developments in Very Large Scale Integration (VLSI) circuit technology have made it possible to fabricate large numbers of very simple processors on a single chip. As most of the processors contained in a shuffle-exchange computer are very simple, the shuffle-exchange graph serves as an excellent basis upon which to design and build chip-sized microcomputers. One of the main difficulties with such an architecture, however, is the problem of routing the wires which link the processors together in a shuffle-exchange network. Current fabrication technology limits the designer to two or three layers of insulated wiring on a chip and demands that he make the chip as small in area as possible.

Abstracted, the designer's problem becomes the mathematical question of how to embed the shuffle-exchange graph in the smallest possible two-dimensional grid. Thompson was the first to formalize the question mathematically. In his thesis [T80], he showed that any *layout* (i.e., embedding in a two-dimensional grid) of the N -node shuffle-exchange graph requires at least $\Omega(N^2/\log^2 N)$ area. In addition, he described a layout requiring only $O(N^2/\log^{1/2} N)$ area. Shortly thereafter, Hoey and Leiserson [HL80] described an embedding for the shuffle-exchange graph in the complex plane (which we refer to as the *complex plane diagram*) and showed how the diagram could be used to find an $O(N^2/\log N)$ -area layout for the N -node shuffle-exchange graph.

In this paper, we investigate the algebraic properties of the complex plane diagram in order to find several $O(N^2/\log^{3/2} N)$ -area layouts for the N -node shuffle-exchange graph. In addition to being *asymptotically* superior to previously discovered layouts, the layouts described in this paper are also superior for *small* values of N . In fact, one of these layouts serves as the basis for the more recent work of Leighton and Miller who have described *optimal* layouts for *small* shuffle-exchange graphs in [LM81].

Subsequent to the completion of the research presented in this paper, we learned that Rodch and Steinberg independently discovered an $O(N^2/\log^{3/2}N)$ -area layout for the N -node shuffle-exchange graph. Their work is also based on the complex plane diagram and appears in [SR81]. Even more recently, Kleitman, Leighton, Lepley and Miller [KLLM81] have discovered an entirely new method for laying out shuffle-exchange graphs which can be used to find *asymptotically optimal* $O(N^2/\log^2N)$ -area layouts. Although their layouts are not entirely practical, they are the only layouts known to achieve Thompson's lower bound asymptotically.

The remainder of the paper is divided into six sections. In section 2, we define the shuffle-exchange graph and the grid model of a chip. We also describe Thompson's $O(N^2/\log^{1/2}N)$ -area layout for the N -node shuffle-exchange graph. In section 3, we define the complex plane diagram for the shuffle-exchange graph and mention several of its properties. In section 4, we describe several layouts for the shuffle-exchange graph which are based on the complex plane diagram. These include a straightforward $O(N^2/\log N)$ -area layout and several new $O(N^2/\log^{3/2}N)$ -area layouts. Section 5 contains some remarks and open questions, and sections 6 and 7 contain the acknowledgements and references.

2. Preliminaries

2a) 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 $w' = a_{k-2}\cdots a_0a_{k-1}$ or $w' = a_0\cdots a_{k-1}a_1$, 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 $w = a_{k-1}\cdots a_10$ and $w' = a_{k-1}\cdots a_11$ or vice-versa). As an example, we have drawn the 8-node shuffle-exchange graph in Figure 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 paper.




Figure 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 is well known to magicians and mathematicians [DGK81] as well as to computer scientists [S71, S80]. For a good survey of the computational power of the shuffle-exchange graph, we recommend Schwartz' paper on ultracomputers [S80]. In addition, Stone's paper [S71] contains a nice description of some important parallel algorithms based on the shuffle-exchange graph.

2b) The 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* [T79, T80]. 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. Wires are routed through the tracks in order to connect pairs of processors. Although a wire in a horizontal track is allowed to cross a wire in a vertical track (without making an electrical connection), pairs of wires are not allowed to overlap for any distance or to overlap at corners (i.e., in they cannot overlap in the same track). Further, wires are not allowed to overlap processors to which they are not linked. (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 2. As before, the shuffle edges are drawn with solid lines while the exchange edges are drawn with dashed lines. Notice that we have omitted the self-loops in Figure 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 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.

2c) Thompson's layout

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 a group, 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 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 [BO78].) 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 $O(N^2/\log^{1/2}N)$.

3. The Complex Plane Diagram

In [HL80], 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 (which we call the *complex plane diagram*) and point out some of its more important properties.

3a) Definition

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

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

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 3. 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.)




Figure 3: The complex plane diagram for the 32-node shuffle-exchange graph. (Taken from [HL80].)

3b) Properties

Examination of Figure 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 is composed of 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 proposition, the number of degenerate necklaces is quite small compared to the number of full necklaces.

Proposition 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

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

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 by 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 3. We will use the properties of levels to find $O(N^2/\log^{3/2}N)$ -area layouts for the N -node shuffle-exchange graph.

4. Layouts Based on the Complex Plane Diagram

In this section, we present several layouts of the shuffle-exchange graph which are based on the complex plane diagram. We commence with a straightforward $O(N^2/\log N)$ -area layout of the N -node shuffle-exchange graph. This layout has been discovered by many researchers (including Hoey and Leiserson). Later, we show how the layout can be modified so as to require only $O(N^2/\log^{3/2}N)$ area.

4a) A straightforward $O(N^2/\log N)$ -area layout

In what follows, we describe a straightforward layout of the shuffle-exchange graph which requires only $O(N^2/\log N)$ area. The layout is formed from a grid of levels and necklaces which we refer to as the *level-necklace grid*. Each row of the grid corresponds to a level of the complex plane diagram. The columns of the grid are divided into consecutive column pairs, each pair corresponding to a necklace. The leftmost column of each column pair corresponds to that part of the necklace which is contained in the left half of the complex plane. Similarly, the rightmost column of each pair corresponds to the part of the necklace contained in the right half of the complex plane.

The rows of the level-necklace grid must have the same top-to-bottom order as do the corresponding levels in the complex plane diagram. The columns, however, may be arranged arbitrarily (provided that columns corresponding to the same necklace are adjacent in the grid).

Each node of the shuffle-exchange graph is placed at the intersection of the row and column of the grid which correspond to the level and part of the necklace (left half or right half) to which it belongs in the complex plane diagram. For example, we have done this for a random ordering of the necklaces of the 32-node shuffle-exchange graph in Figure 4. (Notice that we have used just one column each for the degenerate necklaces $\langle 0 \rangle$ and $\langle 31 \rangle$ since they each contain just one node. In general two columns will be required for necklaces which are mapped to the origin of the complex plane diagram, but the nodes of each such necklace should still be lumped together at a single point of the level-necklace grid.)

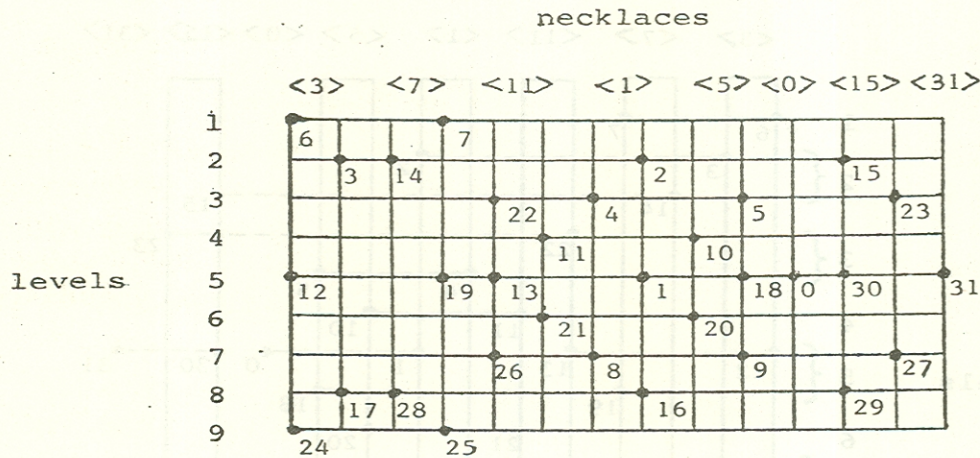


Figure 4: A level-necklace grid for the 32-node shuffle-exchange graph.

Given a level-necklace grid for a shuffle-exchange graph, it is not difficult to produce a layout for the graph. The main step is to partition the exchange edges in each row of the grid into nonoverlapping subsets. Each subset can then be assigned to a horizontal track of the layout. Except for the row corresponding to the real line in the complex plane diagram, the assignment of subsets to horizontal

tracks within a row is arbitrary. (The assignment of horizontal tracks containing nodes on the real line must preserve the cyclic orientation of the nodes which are in necklaces that are mapped to the origin.)

Once this is done, the exchange edges can be inserted in the horizontal tracks and the shuffle edges can be inserted in the vertical tracks. (To be precise, some of the shuffle edges also occupy part of a horizontal track at the top or bottom of the layout.) By Proposition 1, the number of vertical tracks occupied by the necklaces is at most $2N/\log N + O(N^{1/2})$. Since there are precisely $N/2$ exchange edges, at most $N/2 + 2$ horizontal tracks are contained in the layout. Thus the total area of the layout of the N -node shuffle-exchange graph is at most $N^2/\log N + O(N^{3/2})$. As an example, we have displayed in Figure 5 a layout of the 32-node shuffle-exchange graph produced from the level-necklace grid in Figure 4.

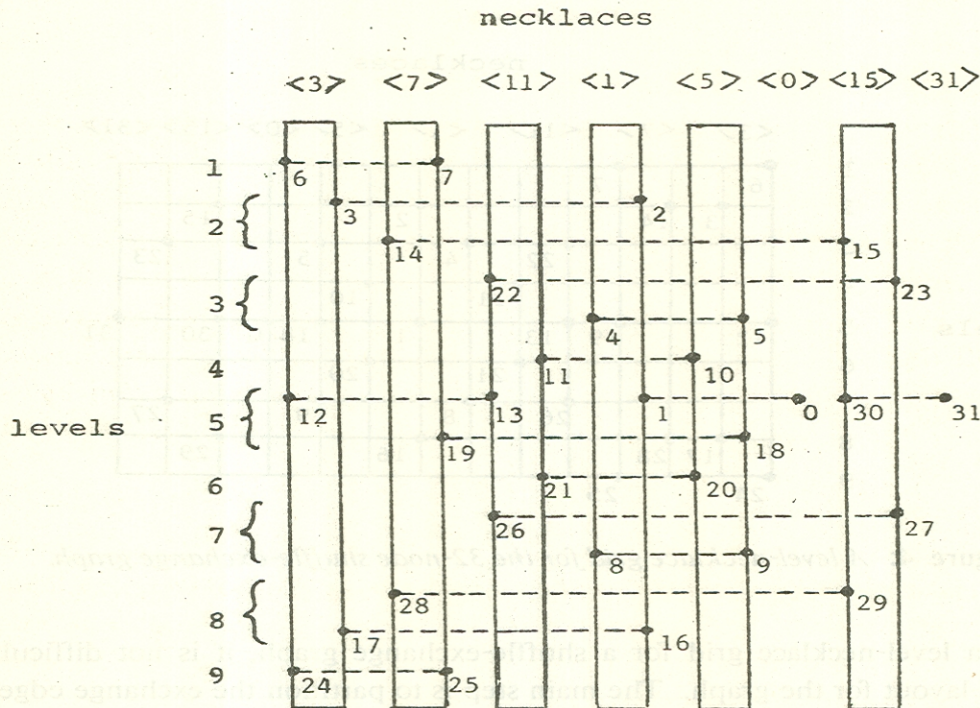


Figure 5: Layout of the 32-node shuffle-exchange graph produced from the level-necklace grid shown in Figure 4.

4b) An improved $O(N^2/\log^{3/2}N)$ -area layout

It is possible to improve the layout described in section 4a by reducing the number of horizontal tracks needed to embed the exchange edges. This can be done by reordering the necklaces from left to right so as to increase the average number of exchange edges which can be inserted on each horizontal track. For example, the ordering of the necklaces shown in Figure 6 results in far fewer horizontal tracks being used than did the ordering of necklaces shown in Figure 5.




Figure 6: An improved layout for the 32-node shuffle-exchange graph.

Although we do not know how to best order the necklaces in general, we have found several orderings which yield $O(N^2/\log^{3/2}N)$ -area layouts for the N -node shuffle-exchange graph. For instance, we will show in what follows that such a layout can be constructed by arranging the necklaces from left to right in order of nondecreasing size. (The *size of a necklace* is simply defined to be the size of any of its nodes.) As an example, the layout displayed in Figure 6 is of this form. (This observation has also been made by Steinberg and Rodeh in [SR81].)

In order to bound the number of horizontal tracks needed to insert the exchange

edges, we will show that the maximum overlap of exchange edges *on each level* is at most the number of nodes of size $h = \lfloor (k-1)/2 \rfloor$ on that level. Since the maximum overlap of exchange edges on each level is an upper bound on the number of horizontal tracks needed to insert the exchange edges on that level, we can thus conclude that the total number of horizontal tracks needed to insert all of the exchange edges is at most

$$B_h \leq B_{k/2} = (2/\pi)^{1/2} N / \log^{1/2} N + O(N / \log^{3/2} N).$$

Thus the resulting layout will have area at most

$$2(2/\pi)^{1/2} N^2 / \log^{3/2} N + O(N^2 / \log^{5/2} N).$$

Although it is clear that the maximum *total* overlap (over all levels) of exchange edges is at most $B_{k/2}$, this is not sufficient to prove the result since any layout must also preserve the top-to-bottom partial order induced by the necklace structure on the exchange edges. It is only within individual levels that the top-to-bottom ordering of exchange edges is arbitrary. (As we noted earlier, some minor precautions are necessary for the level corresponding to the real line.) It is *not* immediately clear, however, why the maximum overlap on each level is at most the number of nodes of size $h \leq k/2$ on that level. In what follows, we establish this result by breaking up each level into sublevels (for which the analysis is easier) and showing that the maximum overlap on each sublevel is at most the number of nodes of size h on that sublevel: The analysis requires some additional notation.

Consider a node of the form $a_{k-1} \cdots a_1 0$ for which either $a_{k-i} = 0$ or $a_i = 0$ or both for each $i \leq k$. We will refer to such a node as *basis node*. A node $b_{k-1} \cdots b_0$ is said to be *generated* by the basis node $a_{k-1} \cdots a_0$ if

- 1) $b_{k-i} = a_{k-i}$ and $b_i = a_i$ whenever $a_{k-i} \neq a_i$ for $1 \leq i \leq k-1$, and
- 2) $b_{k-i} = b_i$ whenever $a_{k-i} = a_i = 0$ for $1 \leq i \leq k-1$.

For example, 10000 generates 10001 , 11100 and 11101 but not 11111 .

It is not difficult to show that if u generates v , then both u and v are on the same level of the complex plane diagram. For example, let $u = a_{k-1} \cdots a_0$ and $v = b_{k-1} \cdots b_0$ and observe that

$$\begin{aligned} p(v) - p(u) &= (b_{k-1} - a_{k-1}) \delta_k^{k-1} + \cdots + (b_1 - a_1) \delta_k + (b_0 - a_0) \\ &= c_{k-1} \delta_k^{k-1} + \cdots + c_1 \delta_k + c_0 \end{aligned}$$

where $c_{k-i} = c_i$ for each i , $1 \leq i \leq k-1$. Since δ_k^{k-i} is the complex conjugate of δ_k^i for $1 \leq i \leq k-1$, we can conclude that $p(v) - p(u)$ is a real number and thus that u and v are in the same level of the complex plane diagram.

It is also easy to show that each node of the shuffle-exchange graph is generated by a unique basis node. In particular, the node which generates $b_{k-1} \cdots b_0$ can be found by

- 1) setting $b_0 = 0$ and (if k is even) setting $b_{k/2} = 0$, and
- 2) setting $b_i = b_{k-i} = 0$ for each i such that (originally) $b_i = b_{k-i} = 1$.

Since exchange edges link nodes which have the same basis node, we can conclude from the preceding arguments that it is possible to partition each level of the complex plane diagram into *sublevels* so that the nodes in each sublevel are precisely the nodes generated by some basis node. We will now show that the maximum overlap on each sublevel is at most the number of nodes of size h on that sublevel.

Since the necklaces have been arranged from left to right in order of nondecreasing size, the overlap of exchange edges between two nodes of size s in any sublevel is at most $O(\max_{0 \leq s \leq k} B_s')$ where B_s' is the number of nodes in that sublevel with size s . In the following proposition, we compute B_s' and show that its maximum for any sublevel occurs at $s = h$.

Proposition 2: *Each basis node of size r generates B_s' nodes of size s , where*

- 1) $B_s' = C(h-r, i)$ for $s = r + 2i$ and $i \leq h-r$, and
- 2) $B_s' = C(h-r, i)$ for $s = r + 2i + 1$ and $i \leq h-r$

when k is odd, and

- 1) $B_s' = C(h-r+1, i)$ for $s = r + 2i$ and $i \leq h-r+1$, and
- 2) $B_s' = 2C(h-r, i)$ for $s = r + 2i + 1$ and $i \leq h-r$

when k is even.

Proof: When k is odd, there are precisely $h-r$ pairs $a_j = a_{k-j} = 0$ in a basis node of size r . In order to generate a string of size $s = r + 2i$ when k is odd, we must set $b_0 = 0$ and set i of the $h-r$ pairs so that $b_j = b_{k-j} = 1$. There are $C(h-r, i)$ such strings. To generate a string of size $s = r + 2i + 1$ when k is odd, we must set $b_0 = 1$ and choose i of the $h-r$ pairs so that $b_j = b_{k-j} = 1$. As before, there are $C(h-r, i)$ such strings.

When k is even, there is also the degenerate pair $a_{k/2} = 0$. To generate a string of size $s = r + 2i$ when k is even, we must choose i of the $h - r + 1$ pairs so that $b_j = b_{k-j} = 1$ (this count includes the "pair" $b_0 = b_{k/2} = 1$). There are $C(h - r + 1, i)$ such strings. To generate a string of size $s = r + 2i + 1$ when k is even, we must set either $b_0 = 1$ and $b_{k/2} = 0$ or $b_0 = 0$ and $b_{k/2} = 1$, and choose i of the $h - r$ pairs so that $b_j = b_{k-j} = 1$ ($j \neq k/2$). There are $2C(h - r, i)$ such strings \square

Given Proposition 2, it is easily checked that the maximum value of B_s' for any sublevel (independent of the value of r) occurs when $s = h$. Thus the sum (over all sublevels) of the maximum overlap at each sublevel is at most the number of nodes of size $h = \lfloor (k-1)/2 \rfloor$ in the entire graph. This is at most $C(k, k/2) \sim (2/\pi)^{1/2} (2^k/k^{1/2})$. Thus the total area of the layout is no more than

$$2(2/\pi)^{1/2} N^2 / \log^{3/2} N + O(N^2 / \log^{5/2} N),$$

as claimed.

4c) Additional $O(N^2 / \log^{3/2} N)$ -area layouts

By varying the order of the necklaces in the level-necklace grid, it is possible to produce a variety of layouts for the shuffle-exchange graph which require at most $O(N^2 / \log^{3/2} N)$ area. The complex plane diagram itself suggests one such ordering. For example, consider an arrangement of the necklaces from left to right in order of nondecreasing radius. (The *radius* of a necklace is defined to be the distance of its nodes from the origin in the complex plane diagram.) Such a layout corresponds to a folding of the complex plane diagram along its imaginary axis followed by a straightening of the necklaces. In what follows, we will show that, like a layout by necklace size, a layout by necklace radius has area $O(N^2 / \log^{3/2} N)$.

Because the layout by radius is so closely related to the complex plane diagram, our analysis will center on the complex plane diagram, itself. As before, we will partition the levels into sublevels and find an upper bound on the maximum overlap of exchange edges on each sublevel separately. The number of horizontal tracks needed to insert the exchange edges will then be at most the sum of these upper bounds. We will show that this sum is at most $O(N / \log^{1/2} N)$.

Notice that the maximum overlap of exchange edges on a sublevel of the level-necklace grid is at most twice the maximum overlap on that sublevel in the complex plane diagram. (The factor of two is introduced by the "folding" of the

diagram along its imaginary axis. Although straightening the necklaces might affect the maximum *total* overlap of exchange edges, it does *not* affect the overlap *within* a sublevel.)

Within a sublevel, an exchange edge can be identified by the real part of its midpoint. For example, the real part of the midpoint of exchange edge $(b_{k-1} \dots b_1 0, b_{k-1} \dots b_1 1)$ is

$$b_{k-1} \cos[2\pi(k-1)/k] + \dots + b_1 \cos[2\pi/k] + 1/2.$$

If a is a basis element of a sublevel, then a generates the other nodes in that sublevel by substitution of the appropriate pairs of ones. For instance, we may set $b_i = b_{k-i} = 1$, if $a_i = a_{k-i} = 0$. Let

$$T_a = \{ 1 \leq j \leq h \mid a_j = a_{k-j} = 0 \}$$

denote those indices $1 \leq i \leq h$ where a pair of 1-bits may be substituted for a pair of 0-bits. (As before, $h = \lfloor (k-1)/2 \rfloor$ but for convenience, we shall henceforth assume that k is odd.) Notice that if b is generated by a , then the real part of the midpoint of the exchange edge incident to b is

$$\sum_{i \in T_a} 2b_i \cos(2\pi i/k) + \sum_{i \notin T_a, 1 \leq i \leq h} \cos(2\pi i/k) + 1/2$$

We now introduce a random variable Z_a , which has as its image, all of the real parts of the midpoints of edges in the sublevel generated by a . Since $b_i = b_{k-i}$ can be either 0 or 1 when $i \in T_a$, let B_i be a random variable representing this choice. In particular,

$$\begin{aligned} B_i &= 0 \quad \text{with probability } 1/2, \quad \text{and} \\ B_i &= 1 \quad \text{with probability } 1/2. \end{aligned}$$

Then

$$\begin{aligned} Z_a &= \sum_{i \in T_a} 2 \cos(2\pi i/k) B_i + \sum_{i \notin T_a, 1 \leq i \leq h} \cos(2\pi i/k) + 1/2 \\ &= \sum_{i \in T_a} 2 \cos(2\pi i/k) (B_i - 1/2). \end{aligned}$$

Since the exchange edges have unit length in the complex plane diagram, two edges overlap if and only if their midpoints are within unit distance of each other. Thus the number of edges which overlap at position x on the sublevel generated by a node a is given by the formula

$$2^{|T_a|} \text{Prob}[x - 1/2 \leq Z_a \leq x + 1/2],$$

where $|T_a|$ denotes the cardinality of T_a . (We caution the reader that the notation $|x|$ is also used to denote the *absolute value* of x .)

Although the distribution function of Z_a is difficult to analyze directly, it does behave like a normal distribution. This is because Z_a is the sum of independent random variables which have mean 0 and variance $\sigma_i^2 = \cos^2(2\pi i/k)$. The Berry-Esseen Theorem states precisely how far Z_a can vary from a normal distribution. (For a proof of this theorem see [F71].)

Berry-Esseen Theorem: Let X_1, X_2, \dots, X_m be independent random variables such that $E(X_i) = 0$, $E(X_i^2) = \sigma_i^2$, and $E(|X_i^3|) = \rho_i$ for $1 \leq i \leq m$. Set $s^2 = \sigma_1^2 + \dots + \sigma_m^2$ and $r = \rho_1 + \dots + \rho_m$. In addition, let F denote the cumulative distribution function of the sum $(X_1 + \dots + X_m)/s$. Then for all x ,

$$|F(x) - \Phi(x)| \leq 6r/s^3$$

where Φ is the standard normal cumulative distribution function \square

In the case of a sublevel generated by a node a , we have

$$X_i = 2 \cos(2\pi i/k) (B_i - 1/2) \quad \text{for } i \in T_a,$$

$$s_a^2 = \sum_{i \in T_a} \cos^2(2\pi i/k) \quad \text{and}$$

$$r_a = \sum_{i \in T_a} |\cos^3(2\pi i/k)|.$$

Applying the Berry-Esseen Theorem, we can thus conclude that

$$\begin{aligned} \text{Prob}[x - 1/2 \leq Z_a \leq x + 1/2] &= \text{Prob}[(x - 1/2)/s_a \leq Z_a/s_a \leq (x + 1/2)/s_a] \\ &\leq \Phi[(x + 1/2)/s_a] - \Phi[(x - 1/2)/s_a] + 12r_a/s_a^3 \end{aligned}$$

Because the standard normal density function is symmetric and unimodal, we can conclude that the maximum of $\text{Prob}[x - 1/2 \leq Z_a \leq x + 1/2]$ occurs at $x = 0$ and is at most $O(1/s_a + r_a/s_a^3)$.

In the following proposition, we find bounds for the values of r_a and s_a .

Proposition 3: For any basis node a

$$r_a = \sum_{i \in T_a} |\cos^3(2\pi i/k)| \leq |T_a| \quad \text{and}$$

$$s_a^2 = \sum_{i \in T_a} \cos^2(2\pi i/k) \geq \Omega(|T_a|^3/k^2).$$

Proof: The bound on r_a is easy to compute since $|\cos^3(2\pi i/k)| \leq 1$. The calculation of s_a is a bit more tedious. In order to obtain a lower bound, $\cos^2(2\pi i/k)$ must be made as small as possible. The smallest values occur when T_a contains indices i which are as close to $(k-1)/4$ as possible. In this case, we can approximate $\cos^2(2\pi i/k)$ with the value $c(\pi/2 - 2\pi i/k)^2$, for some constant c . Direct computation reveals that the sum of these squares is at least $\Omega(|T_a|^3/k^2)$ \square

Since $|T_a| < k$ for all a , we can conclude from the preceding that the maximum overlap of exchange edges on a sublevel generated by a is at most

$$O(2^{|T_a|} k^3 / |T_a|^{7/2}).$$

Noting that there are precisely $C(h, j) 2^{h-j}$ sublevels generated by a node for which $|T_a| = j$ and summing, we can conclude that the total number of horizontal tracks needed to insert all of the exchange edges is at most

$$\sum_{j=1}^h C(h, j) 2^{h-j} O(2^j k^3 / j^{7/2})$$

$$= O[k^3 2^h \sum_{j=1}^h C(h, j) / j^{7/2}].$$

It is not difficult to check that the dominant terms in the preceding sum occur when $j = h/2 \pm \Theta(h^{1/2} \log h)$. In this region, $j = \Theta(k)$ and thus the sum is bounded above by

$$O[2^h k^{1/2} \sum_{j=1}^h C(h, j)] = O(2^{k-1} / k^{1/2})$$

$$= O(N / \log^{1/2} N),$$

thus completing the proof that a layout by necklace radius takes at most $O(N^2 / \log^{3/2} N)$ area.

5. Remarks

It is worth remarking that the $O(N^2/\log^{3/2}N)$ -area layouts for the shuffle-exchange graph described in section 4 actually require $\Omega(N^2/\log^{3/2}N)$ area and thus our analysis of these layouts cannot be improved by more than a constant factor. In each case, the lower bound on area can be derived from the fact that the maximum *total* overlap of exchange edges in the layouts is at least $\Omega(N/\log^{1/2}N)$. (Remember that although the maximum *total* overlap of exchange edges is *not* an *upper* bound on the number of horizontal tracks needed to insert the exchange edges, it *is* a lower bound.)

The $\Omega(N/\log^{1/2}N)$ lower bound on maximum overlap is easily established for the layout according to necklace size since $\Omega(N/\log^{1/2}N)$ exchange edges link nodes of size $k/2$ to nodes of size $k/2 + 1$. The lower bound on maximum overlap is somewhat more difficult to prove for the layout according to necklace radius. The first step in the proof is to show that at least $N/2$ exchange edges are contained within a square of side length $ck^{1/2}$ centered at the origin of the complex plane diagram (where c is a constant). (This can be done by using the techniques developed in section 4c.) Next consider the sum (over i) of the total overlaps at points corresponding to radii of $i/2$ for $1 \leq i \leq ck^{1/2}$. Because the complex plane diagram is radially symmetric, it is possible to show that at least $\Omega(N)$ exchange edges are counted in this sum. Thus the overlap at one of these points must be at least $\Omega(N/k^{1/2})$.

Since Thompson [T80] has shown that any layout for the N -node shuffle-exchange graph must have area at least $\Omega(N^2/\log^2N)$, we know that at least $\Omega(N/\log N)$ horizontal tracks are needed to insert the exchange edges for any ordering of necklaces in the level-necklace grid. However, there is no ordering of the necklaces known for which the exchange edges can be inserted using less than $o(N/\log^{1/2}N)$ horizontal tracks. This suggests an interesting open question since it would be nice to find an $O(N^2/\log^2N)$ -area layout based on the complex plane diagram. (Although an asymptotically optimal $O(N^2/\log^2N)$ -area layout for the shuffle-exchange graph has recently been found by Kleitman, Leighton, Lepley and Miller [KLLM81], it is rather complicated and of limited practical use.)

Although we do not know of necklace orderings for which the exchange edges can be inserted using less than $o(N/\log^{1/2}N)$ horizontal tracks, we *do* know of orderings for which the *maximum total overlap* of exchange edges is at most

$O(N \log \log N / \log N)$. For example, an ordering of the necklaces by minimum value has a maximum total overlap of $\Theta(N \log \log N / \log N)$. (The *minimum value* of a necklace is simply the minimum of the values of the nodes contained in the necklace.)

Interestingly, an analysis of the minimum (over all orderings) of the maximum total overlap for small values of N indicates that there may always be an ordering for which the maximum total overlap is at most $O(N / \log N)$, the least possible. In fact, for $3 \leq N \leq 7$, this minimum maximum overlap is precisely $\lfloor (2^k - 2) / k \rfloor$. A summary of the minimum maximum overlap data for small values of N is included in Table 1.

Table 1

Maximum Overlap of Best Known Orderings

k	N	maximum overlap of best known ordering	optimal?
3	8	2	yes
4	16	3	yes
5	32	6	yes
6	64	10	yes
7	128	18	yes
8	256	33	yes
9	512	62	?
10	1024	115	?
11	2048	214	?
12	4096	388	?
13	8192	754	?

In addition to varying the order of the necklaces, improvements in the layout may also be made by rearranging the level assignments of the exchange edges. For example, the layout of the 32-node shuffle-exchange graph shown in Figure 7 was constructed in this way. (The careful reader will notice that we have also

manipulated the necklaces somewhat in order to produce this layout.) For a more detailed discussion of the manner in which exchange edges can be reassigned, we refer the reader to [LM81]. (Such layouts have also been used in conjunction with the Blue Chip Project at Purdue [S81].)




Figure 7: An improved layout for the 32-node shuffle-exchange graph.

6. Acknowledgements

In acknowledgement, we would like to thank the following people for their helpful remarks and suggestions: Herman Chernoff, Peter Elias, Dan Hoey, Dan Kleitman, Charles Leiserson, Ron Rivest, Michael Rodeh, Larry Snyder, and Richard Zippel.

7. References

- [BO78] C. M. Bender and S. A. Orszag, *Advanced Mathematical Methods for Scientists and Engineers*, McGraw-Hill Book Company, New York, 1978.
- [DGK81] P. Diaconis, R. L. Graham and W. M. Kantor, "The mathematics of perfect shuffles," preprint, 1981.
- [F71] W. Feller, *An Introduction to Probability Theory and Its Applications*, Volume II, second ed., John Wiley and Sons, New York, 1971.
- [HL80] D. Hoey and C. E. Leiserson, A layout for the shuffle-exchange network," *Proceedings of the 1980 IEEE International Conference on Parallel Processing*, August 1980.

- [KLLM81] D. Kleitman, F. T. Leighton, M. Lepley and G. L. Miller, "New layouts for the shuffle-exchange graph," *Proceedings of the 13th Annual ACM Symposium on Theory of Computing*, May 1981, pp. 278-292.
- [L81] F. T. Leighton, *Layouts for the Shuffle-Exchange Graph and Lower Bound Techniques for VLSI*, Ph.D. Thesis, Mathematics Department, Massachusetts Institute of Technology, Cambridge Massachusetts, September 1981. (Also to appear as an MIT Technical Report.)
- [LM81] F. T. Leighton and G. L. Miller, "Optimal layouts for small shuffle-exchange graphs," *VLSI 81 - Very Large Scale Integration*, edited by John P. Gray, Academic Press, London, August 1981, pp. 289-299.
- [P80] D. S. Parker, "Notes on shuffle/exchange-type switching networks," *IEEE Transactions on Computers*, Vol. C-29, No. 3, March 1980, pp. 213-222.
- [S71] H. S. Stone, "Parallel processing with the perfect shuffle," *IEEE Transactions on Computers*, Vol. C-20, No. 2, February 1971, pp. 153-161.
- [S80] J. T. Schwartz, "Ultracomputers," *ACM Transactions on Programming Languages and Systems*, Vol. 2, No. 4, October 1980, pp. 484-521.
- [S81] L. Snyder, "Overview of the CHiP computer," *VLSI 81 - Very Large Scale Integration*, edited by J. Gray, Academic Press, London, August 1981, pp. 237-246.
- [SR81] D. Steinberg and M. Rodeh, "A layout for the shuffle-exchange network with $\Theta(N^2/\log^{3/2}N)$ area," submitted to *IEEE Transactions on Computers*.
- [T79] C. D. Thompson, "Area-time complexity for VLSI," *Proceedings of the 11th Annual ACM Symposium on Theory of Computing*, May 1979, pp. 81-88.
- [T80] C. D. Thompson, *A Complexity Theory for VLSI*, Ph.D. dissertation, Department of Computer Science, Carnegie-Mellon University, 1980.

