An Asymptotically Optimal Layout for the Shuffle–Exchange Graph*

Daniel Kleitman, Frank Thomson Leighton, Margaret Lepley, and Gary L. Miller

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

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 quite simple and many require no more than logarithmic time and constant space per processor. An $O(N^2/\log^2 N)$-area layout for the shuffle–exchange graph on a two-dimensional grid is described. The layout is the first which is known to achieve Thompson's asymptotic lower bound.

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 [15, 18, 19]. The algorithms needed for these operations are very 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 simple processors on a single chip. As most of the processors contained in a shuffle–exchange computer are 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 [23]

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 [4] described an embedding for the suffle–exchange graph in the complex plane (called 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. Subsequently, Leighton *et al.* [12], and (independently) Steinberg and Rodeh [21] showed how the complex plane diagram could be used to find an $O(N^2/\log^{3/2}N)$-area layout.

In this paper, we pursue an entirely different strategy in order to find an $O(N^2/\log^2 N)$-area layout for the N-node shuffle–exchange graph, thus achieving Thompson's asymptotic lower bound. In addition, we describe how the new techniques can be used to find $O(N^2/\log^2 N)$-area layouts for more general graphs (such as the shuffle–shift–reverse graph).

As is the case with much of the previous work on this problem, our results are more theoretical than practical. This is due, in part, to the fact that the layout procedure described in the paper is designed to produce good layouts for *large* shuffle–exchange graphs. Unfortunately, it produces poor layouts for *small* shuffle–exchange graphs and, for the time being, these are the networks of practical interest. Nevertheless, the *methods* developed in the paper do appear to have some practical value. For example, Leighton and Miller [13] have constructed a 19-by-36 layout for the 128-node shuffle–exchange graph by extending the methods described in this paper.

The remainder of the paper is divided into seven sections. In Section 2, we define the shuffle–exchange graph and the grid model of a chip. Section 3 contains more definitions and some useful combinatorial lemmas. The proofs of these lemmas are included in Section 4. In Section 5, we describe a near-optimal, $O(N^2 \log \log^2 N/\log^2 N)$-area layout for the N-node shuffle–exchange graph. In Section 6, we show how to modify the near-optimal layout in order to produce an asymptotically optimal layout. In Section 7, we show how to lay out several supergraphs of the shuffle–exchange graph. We conclude in Section 8 with some remarks.

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 $w' = a_{k-2} \cdots a_0 a_{k-1}$ or $w' = a_0 a_{k-1} \cdots a_1$, 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 $w = a_{k-1} \cdots a_1 0$ and $w' = a_{k-1} \cdots a_1 1$ or

vice versa). As an example, we have drawn the 8-node shuffle–exchange graph in Fig. 1. Note that the shuffle edges are drawn with solid lines while the exchange edges are drawn with dashed lines.

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 $a_2 a_1 a_0$ sends its card across a shuffle edge to the neighboring processor $a_1 a_0 a_2$. 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 [1] as well as to computer scientists [18, 19]. For a good survey of the computational power of the shuffle–exchange graph, we recommend Schwartz' paper on ultracomputers [19]. In addition, Stone's paper [18] contains a nice description of some important parallel algorithms based on the shuffle–exchange graph.

2b. *Necklaces*

As can be seen in Fig. 1, there is a natural partitioning of the shuffle edges into cycles. These cycles are called *necklaces*. A necklace is simply the collection of all cyclic shifts of some node of the graph. In particular, the necklace which contains the node w is called the *necklace generated by* w and is denoted by $\langle w \rangle$. For example, the necklace generated by 001 is $\langle 001 \rangle = \{001, 010, 100\}$.

If a necklace contains precisely k nodes, it is said to be *full*. Otherwise the necklace contains less than k nodes and is said to be *degenerate*. As Leighton *et al.* show in

FIG. 1.   The 8-node shuffle–exchange graph.

|12| the number of degenerate necklaces is quite small compared to the number of full necklaces. In particular, they prove

LEMMA 1.   *At most* $O(N^{1/2})$ *nodes are contained in the degenerate necklaces of an N-node shuffle-exchange graph. The remaining nodes are contained in* $N/\log N - O(N^{1/2}/\log N)$ *full necklaces.*

2c. *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* |22, 23|. 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., 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 Fig. 2. Notice that we have omitted the self-loops in Fig. 2 since they are electrically redundant. In general, the embedding will not be planar (as it is in this example).

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 Fig. 2 has area 18. As Leighton and Miller |13| have shown, this layout is suboptimal.

Leiserson |7| observed that any M wires can be added to a layout by inserting at most $2M$ vertical and $2M$ horizontal tracks. Hence M wires can be added to a $\Omega(M)$-by-$\Omega(M)$ layout without increasing the area by more than a constant factor. As any
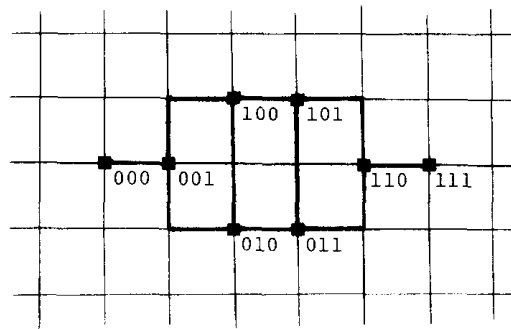

FIG. 2.   A grid model layout of the 8-node shuffle-exchange graph.

layout for the N-node shuffle–exchange graph must have $\Omega(N/\log N)$ vertical and $\Omega(N/\log N)$ horizontal tracks, the preceding observation means that a nearly complete layout for the shuffle–exchange graph with area A can be extended to a complete layout with area $O(A)$. This result will be used at several points in the paper and is stated formally in

LEMMA 2. *Any area A layout which contains all but $O(N/\log N)$ nodes and edges of the N-node shuffle–exchange graph can be extended to form a complete layout for the N-node shuffle–exchange graph with area $O(A)$.*

As an immediate application of Lemma 2, we can henceforth ignore nodes which are contained in degenerate necklaces. This is because at most $O(N^{1/2})$ nodes are contained in degenerate necklaces (Lemma 1) and thus they can be inserted into any layout of full necklaces without increasing the total area by more than a constant factor (Lemma 2).

3. SOME COMBINATORIAL LEMMAS

In what follows, we will be particularly interested in the size and location of the longest block of consecutive 0-bits in the k-bit binary string associated with each node. In order that the size of this block be the same for all nodes within a necklace, we allow blocks to begin at the end and end at the beginning of a string. For example, the longest block of zeros in the string 01010 starts at the fifth bit and has length two.

Let $\Psi_k(t)$ denote the number of k-bit strings for which the longest block of consecutive zeros has length t. For example, $\Psi_3(2) = 3$. The following combinatorial lemma provides a good asymptotic bound on the growth of $\Psi_k(t)$. The proof of the lemma (as well as of Lemmas 4–6) is included in Section 4.

LEMMA 3. *For $(\log k)/2 + \log \ln k \leqslant t \ll k$ and $k \to \infty$,*

$$\Psi_k(t) \sim 2^k (e^{-k2^{-(t+2)}} - e^{-k2^{-(t+1)}}).$$

In order to illustrate the important features of the function in Lemma 3, we have sketched a graph of $2^{-k}\Psi_k(t)$ versus t in Fig. 3. The maximum of $2^{-k}\Psi_k(t)$ occurs at $t = \log k - 1$, where $2^{-k}\Psi_k(t) \sim (e^{1/2} - 1)/e \approx 0.23865$. For $t > \log k - 1$, $2^{-k}\Psi_k(t)$ decreases *exponentially* as t increases. For $t \leqslant \log k - 1$, $2^{-k}\Psi_k(t)$ decreases *doubly exponentially* as t decreases.

Roughly speaking, Lemma 3 states that the longest block of consecutive zeros in nearly $\frac{1}{4}$ of all k-bit strings has length precisely $\log k - 1$. Further, there are not many strings of length k with substantially more than $\log k$ consecutive zeros and even fewer strings for which the longest block of consecutive zeros has length substantially less than $\log k$. This information is further quantified in Lemma 4:
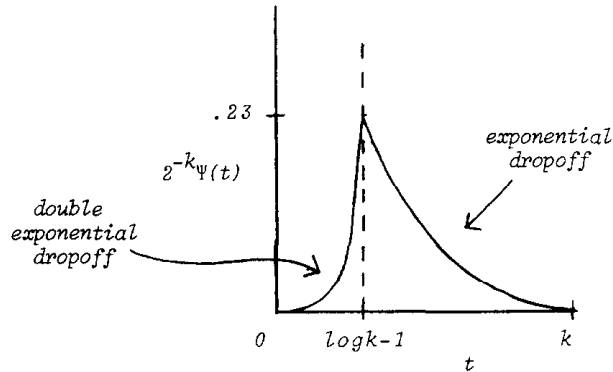
FIG. 3. Density of k-bit binary strings for which the longest block of consecutive zeros has length t.

LEMMA 4. *The number of k-bit strings for which the longest block of consecutive zeros has length less than* $\log k - \log \ln k - 1$ *or length greater than* $2 \log k$ *is at most* $O(2^k/k) = O(N/\log N)$.

By Lemma 2, we may ignore $O(N/\log N)$-sized sets of nodes which have undesirable properties. As such nodes can be inserted with the addition of at most $O(N/\log N)$ vertical and horizontal tracks, we can always add them later without increasing the total area by more than a constant factor. By Lemma 4, we can thus henceforth consider only those nodes for which the longest block of zeros has length between $\log k - \log \ln k - 1$ and $2 \log k$.

We will also be interested in the size of the *second* longest block of consecutive zeros in each string. Usually, the size of the second longest block of zeros will be very close to the size of the longest block of zeros. We state this observation more precisely in

LEMMA 5. *The sum over all necklaces of the difference in length between the longest and second longest blocks of consecutive zeros is at most* $O(N/\log N)$.

Using information about the size and location of blocks of zeros within the necklace, it is possible to distinguish one particular node in most necklaces. More precisely, we define the *distinguished node of a necklace* to be the node containing the longest leading block of zeros. For example, 00101 is the distinguished node of $\langle 01010 \rangle$. Should two or more nodes of a necklace begin with equal and maximum length blocks of zeros, then each node of the necklace contains at least two blocks of zeros of maximum length. In such cases, we distinguish that node for which the leading block of zeros has maximum length and for which the second occurrence of a *maximum* length block of zeros is as near as possible to the beginning of the string. For example, 01011 (not 01101) is the distinguished node of the necklace $\langle 10101 \rangle$. For some necklaces, such as $\langle 111 \rangle$ and $\langle 1010101 \rangle$, there is no uniquely distinguished node. As we show in Lemma 6 such necklaces are sufficiently rare that we need not consider them further.

LEMMA 6. *At most $O(N/\log N)$ nodes are contained in necklaces which fail to have a uniquely distinguished node.*

We refer to the leading block of zeros of a distinguished node as the *primary block of zeros*. If a distinguished node has two or more maximum length blocks of zeros, then the maximum length block following the primary block is referred to as the *secondary block of zeros*. These definitions can be easily extended to any node contained in a necklace which has a uniquely distinguished node. For example, the primary block of zeros of 01010 starts in the fifth bit and has length two. Note that this string does *not* have a secondary block of zeros. As another example, we note that the secondary block of zeros in the string 11010 consists solely of the fifth bit. Note that the secondary block of zeros (if it exists) always has the same length as the primary block of zeros.

If the last bit of a node occurs in the primary block of zeros, we call that node a *primary node*. Similarly, if the last bit of a node occurs in the secondary block of zeros, we call the node a *secondary node*. For example, 10110 is a primary node, 11010 is a secondary node and 10010 is neither primary nor secondary.

Note that all primary and secondary nodes are necessarily even. (We say that a node is *even* if its last bit is 0 and *odd* if its last bit is 1.) Note also that, by Lemmas 2 and 4, we need only consider necklaces which contain between $\log k - \log \ln k - 1$ and $2 \log k$ primary nodes. Such necklaces will also have at most $2 \log k$ secondary nodes.

In what follows, we will represent nodes in terms of their corresponding distinguished nodes. More precisely, we use the notation $a_{k-1} \cdots a_{i+1} \bar{a}_i a_{i-1} \cdots a_0$ to denote the node $a_{i-1} \cdots a_0 a_{k-1} \cdots a_i$. For example, $00\overline{1}01$ denotes the node 10010. Using this notation, a primary node has the form $0 \cdots \bar{0} \cdots 0w$ while a secondary node has the form $0 \cdots 0w'0 \cdots \bar{0} \cdots 0w''$, where $0 \cdots 0w$ and $0 \cdots 0w'0 \cdots 0w''$ are assumed to be distinguished nodes.

4. PROOFS OF LEMMAS 3–6

We now present the proofs of Lemmas 3–6. Such results can also be found in the recent work of Guibas and Odlyzko [2, 3]. As the proofs are fairly technical, many readers may wish to skip this section and proceed directly to Section 5.

In what follows, we will write $\bar{\Psi}_k(t)$ to denote the number of k-bit strings which do *not* contain $t - 1$ consecutive zeros. Except for the string of all zeros (which we ignore), these are precisely the strings which do not contain the substring

$$v_t = \overbrace{10 \cdots 0}^{t}.$$

The proofs of Lemmas 3–6 depend heavily on the following combinatorial result:

THEOREM 1. *For large t and k,*

$$\overline{\Psi}_k(t) = 2^k e^{-k2^{-t}} e^{O(t2^{-2}, kt2^{-2t})}.$$

Proof. We first count the number $\overline{\Psi}_k'(t)$ of k-bit strings which do not contain an occurrence of v_t between the beginning and end of the string (i.e., for the time being we ignore the occurrences of v_t which begin at the end and end at the beginning of a string).

Fix t and let f_i denote the number of i-bit strings ending with v_t which do not contain any other occurrences of v_t in the string. Set $F(x) = \sum_{i=0}^{\infty} f_i x^i$. Note that $\overline{\Psi}_k'(t)$ is the $(k + t)$th coefficient of $F(x)$. Let $f_i^{(j)}$ denote the number of i-bit strings ending in v_t which contain precisely j occurrences of v_t and set

$$F^{(j)}(x) = \sum_{i=0}^{\infty} f_i^{(j)} x^i.$$

Since occurrences of v_t cannot overlap, it is not difficult to show that $F^{(j)}(x)$ is identical to $F(x)^j$ for all $j > 1$.

Let g_i be the number of i-bit strings which end in v_t (regardless of the number of other occurrences of v_t which appear in the string) and set $G(x) = \sum_{i=0}^{\infty} g_i x^i$. Since $g_i = 2^{i-t}$ for all $i \geq t$, it is easily seen that $G(x) = x^t/(1 - 2x)$. Also note that

$$G(x) = \sum_{j=1}^{\infty} F^{(j)}(x) = \sum_{j=1}^{\infty} F(x)^j = [1/(1 - F(x))] - 1$$

and thus that

$$F(x) = G(x)/(G(x) + 1) = x^t/(1 - 2x + x^t).$$

Thus $\overline{\Psi}_k'(t)$ is simply the kth coefficient of $1/(1 - 2x + x^t)$. For example, $\overline{\Psi}_4'(2) = 5$ which is the coefficient of x^4 in the expansion of $1/(1 - 2x + x^2)$.

Let $p(x) = 1 - 2x + x^t$. It is easily observed that $\gcd(p(x), dp(x)/dx) = 1$ and thus that $p(x)$ does not have any multiple roots for $t > 2$. Thus we can expand

$$p(x)^{-1} = \sum_{i=1}^{t} A_i/(x - r_i),$$

where $\{r_i \mid 1 \leq i \leq t\}$ is the set of distinct (and possibly complex) roots of $p(x)$ and

$$A_i = [(x - r_i)/p(x)]_{x=r_i} = 1/[dp(x)/dx]_{x=r_i}$$

for $1 \leq i \leq t$. Once the roots of $p(x)$ are known, we can calculate $\overline{\Psi}_k'(t)$ from the formula

$$\overline{\Psi}_k'(t) = - \sum_{i=1}^{t} A_i r_i^{-(k+1)}.$$

Although we do not know how to find the roots of $p(x)$ explicitly for large t, we can describe them asymptotically. First observe that as $t \to \infty$, the absolute value of every root must approach either $\frac{1}{2}$ or 1. Otherwise the absolute value of one term of $p(x)$ would dominate the sum of the absolute values of the other two terms. For example, if $|r| \leqslant c < \frac{1}{2}$ as $t \to \infty$ for some root r and constant c, then $1 > |2r| + |r^t|$ for large t.

If there are to be any roots r such that $|r| \to \frac{1}{2}$, it is essential that $r \to \frac{1}{2}$. Otherwise, the real part of $p(r)$ cannot vanish for large t. By substituting $\frac{1}{2} e^{s(t)}$ for r, where $s(t) \to 0$ as $t \to \infty$, we find that

$$1 - e^{s(t)} + 2^{-t} e^{ts(t)} = 0$$

and thus that

$$1 - (1 + s(t) + O(s(t)^2)) + 2^{-t}(1 + O(ts(t))) = 0.$$

Thus $s(t) = 2^{-t} + q(t)$, where $|q(t)| \ll 2^{-t}$ as $t \to \infty$. Another iteration of this process reveals that $q(t) = O(t2^{-2t})$ and thus that

$$r = \tfrac{1}{2} e^{2^{-t}} e^{O(t2^{-2t})} \qquad \text{as} \quad t \to \infty.$$

In fact, there is precisely one root, say r_1, which approaches $\frac{1}{2}$ as $t \to \infty$. The absolute values of the remaining roots approach 1. In particular, the absolute values of these roots must be greater than or equal to 1 for large t. Otherwise there would be a root r and a function $\varepsilon(t) \to 0^+$ such that $|r| = 1 - \varepsilon(t)$. But then

$$|2r| = 2 - 2\varepsilon(t) > 1 + |1 - \varepsilon(t)|^t = 1 + |r^t|$$

for $t > 2$ and it would be impossible for $p(r)$ to vanish for large t, a contradiction.

It remains to compute the A_i. Since $dp(x)/dx = tx^{t-1} - 2$, we find that $A_1 = -\frac{1}{2} + O(t2^{-t})$ and that $A_i = O(1/t)$ for $2 \leqslant i \leqslant t$. Thus

$$\bar{\Psi}'_k(t) = O(1) - [-\tfrac{1}{2} + O(t2^{-t})] \, 2^{k+1} e^{-(k+1)2^{-t}} e^{O(kt2^{-2t})}.$$

Replacing $1 + O(t2^{-t})$ with $e^{O(t2^{-t})}$ and simplifying, we conclude that

$$\bar{\Psi}'_k(t) = 2^k e^{-k2^{-t}} e^{O(t2^{-t}, kt2^{-2t})}$$

for large t and k.

The only strings which are included in the count of $\bar{\Psi}'_k(t)$ but not in that of $\bar{\Psi}_k(t)$ are those of the form

$$\overbrace{0 \cdots 0}^{i} w \overbrace{10 \cdots 0}^{t-i},$$

where $1 \leqslant i \leqslant t - 1$ and w is a string which is included in the count of $\bar{\Psi}'_{k-i}(t)$. Thus

$$\bar{\Psi}_k(t) = \bar{\Psi}'_k(t) - (t-1)\,\bar{\Psi}'_{k-i}(t)$$

$$= 2^k e^{-k2^{-t}} e^{O(t2^{-t},\,kt2^{-2t})} - (t-1)\,2^{k-t} e^{-(k-t)2^{-t}} e^{O(t2^{-t},\,kt2^{-2t})}$$

$$= 2^k e^{-k2^{-t}} e^{O(t2^{-t},\,kt2^{-2t})}$$

for large t and k. This completes the proof of the theorem. ∎

We can now prove Lemmas 3 and 4.

Proof of Lemma 3.　From the definition, we know that

$$\Psi_k(t) = \bar{\Psi}_k(t+2) - \bar{\Psi}_k(t+1)$$

$$= 2^k e^{-k2^{-(t+2)}} e^{O(t2^{-t},\,kt2^{-2t})} - 2^k e^{-k2^{-(t+1)}} e^{O(t2^{-t},\,kt2^{-2t})}$$

for large t and k. For $t \geqslant (\log k)/2 + \log \log k$, both $t2^{-t}$ and $kt2^{-2t}$ vanish as $k \to \infty$. In what follows, we will show that if $t \ll k$, then

$$e^{-k2^{-(t+2)}} - e^{-k2^{-(t+1)}} \gg O(t2^{-t},\,kt2^{-2t})$$

and thus that

$$\Psi_k(t) \sim 2^k (e^{-k2^{-(t+2)}} - e^{-k2^{-(t+1)}}).$$

Assume for the purposes of contradiction that

$$e^{-k2^{-(t+2)}} - e^{-k2^{-(t+1)}} \leqslant O(t2^{-t},\,kt2^{-2t}).$$

Then, $e^{-k2^{-(t+2)}} \sim e^{-k2^{-(t+1)}}$ which means that $e^{-k2^{-(t+2)} + k2^{-(t+1)}} \sim 1$ and thus that $k2^{-(t+2)} \to 0$. Thus we can use a Taylor series expansion of the exponentials to find that

$$e^{-k2^{-(t+2)}} - e^{-k2^{-(t+1)}} \sim (1 - k2^{-(t+2)}) - (1 - k2^{-(t+1)})$$

$$= k2^{-(t+2)}$$

$$\gg O(t2^{-t},\,kt2^{-2t})$$

provided that $t \ll k$, a contradiction. ∎

Proof of Lemma 4.　The number of k-bit strings which do not contain a block of $\log k - \log \ln k - 1$ consecutive zeros is

$$\bar{\Psi}_k(\log k - \log \ln k) \sim 2^k e^{-k2^{-\log k + \log \ln k}}$$

$$= 2^k/k$$

$$= O(N/\log N).$$

The number of k-bit strings which contain a block of $2 \log k + 1$ consecutive zeros is

$$2^k - \overline{\Psi}_k(2 \log k + 2) \sim 2^k - 2^k e^{-k2^{-2\log k - 2}} e^{O((\log k)/k^2)}$$

$$= 2^k - 2^k [1 - 1/(4k) + O((\log k)/k^2)]$$

$$\sim 2^k/4k$$

$$= O(N/\log N). \quad \blacksquare$$

The proofs of Lemmas 5 and 6 depend on the following corollary to Theorem 1:

COROLLARY 1.   *For bounded $m > 0$ and p, and large k and t,*

$$\sum_{t=1}^{(k+p)/m} \overline{\Psi}_{k-mt+p}(t) = O(2^k/k^m).$$

Proof.   We first observe that for $t < 2 \log k/3$,

$$\overline{\Psi}_{k-mt+p}(t) \leqslant \overline{\Psi}_k(2 \log k/3) \sim 2^k e^{-k2^{-(2\log k)/3}} = 2^k e^{-k^{1/3}}$$

and thus that

$$\sum_{t=1}^{2\log k/3} \overline{\Psi}_{k-mt+p}(t) \leqslant \tfrac{2}{3} 2^k e^{-k^{1/3}} \log k \ll 2^k/k^m$$

for any finite m and p as $k \to \infty$.

For larger values of t,

$$\overline{\Psi}_{k-mt+p}(t) \sim 2^{k-mt+p} e^{-k2^{-t}}$$

and thus

$$\sum_{t=2\log k/3}^{(k+p)/m} \overline{\Psi}_{k-mt+p}(t) \sim \sum_{t=2\log k/3}^{(k+p)/m} 2^{k-mt+p} e^{-k2^{-t}}.$$

By making the change of variables $r = t - \log k$, we can see that the preceding sum is at most

$$(2^{k+p}/k^m) \sum_{r=-\infty}^{\infty} 2^{-mr} e^{-2^{-r}}$$

and thus at most $O(2^k/k^m) = O(N/\log N). \quad \blacksquare$

Proof of Lemma 5.   A string whose longest block of zeros has length t and whose second longest block of zeros has length $s \leqslant t$ is of the form

$$w1\overbrace{0 \cdots 0}^{t+1}w',$$

where the longest block of zeros in ww' has length s. By definition, there are at most $k\Psi_{k-t-1}(s)$ such strings. Thus the sum over all *necklaces* of the difference between the sizes of the longest block and second longest block of zeros is

$$\leqslant (1/k) \sum_{t=0}^{k-1} \sum_{s=0}^{t} (t-s) k\Psi_{k-t-1}(s)$$

$$= \sum_{t=0}^{k-1} \sum_{s=0}^{t} (t-s)[\bar{\Psi}_{k-t-1}(s+2) - \bar{\Psi}_{k-t-1}(s+1)]$$

$$= \sum_{s=1}^{k} \sum_{t=s}^{k} \bar{\Psi}_{k-t}(s)$$

$$= \sum_{s=1}^{k} \left(2^k e^{-k2^{-s}} e^{O(s2^{-s}, ks2^{-2s})} \sum_{t=s}^{k} 2^{-t} e^{t2^{-s}}\right)$$

$$\leqslant \sum_{s=1}^{k} (2^k e^{-k2^{-s}} e^{O(s2^{-s}, ks2^{-2s})} 2^{-s} e^{O(s2^{-s})})$$

$$= \sum_{s=1}^{k} 2^{k-s} e^{-k2^{-s}} e^{O(s2^{-s}, ks2^{-2s})}$$

$$\leqslant \sum_{s=1}^{k} \bar{\Psi}_{k-s}(s)$$

$$= O(N/\log N)$$

by Corollary 1. ∎

Proof of Lemma 6. Consider a necklace which fails to have a uniquely distinguished node. Each node in such a necklace must have one of the following three forms:

(1) $w_1 \overbrace{0 \cdots 0}^{k/2} w_2 \underbrace{0 \cdots 0}_{t} w_3$,
where the first block is $\underbrace{}_{t}$

(2) $w_1 \overbrace{0 \cdots 0}^{r} w_2 \overbrace{0 \cdots 0}^{r} w_3 0 \cdots 0 w_4$, or
with blocks $\underbrace{}_{t}\underbrace{}_{t}\underbrace{}_{t}$

(3) $w_1 \overbrace{0 \cdots 0}^{r} w_2 0 \cdots 0 w_3 \overbrace{0 \cdots 0}^{r} w_4 0 \cdots 0 w_5$,
with blocks $\underbrace{}_{t}\underbrace{}_{t}\underbrace{}_{t}\underbrace{}_{t}$

where t is the length of the longest block of zeros in any of the strings. It is easily seen that there are at most

(1) $k \sum_{t=1}^{k/2} \overline{\Psi}_{k-2t}(t+2)$ nodes of the first type,

(2) $k^2 \sum_{t=1}^{k/3} \overline{\Psi}_{k-3t}(t+2)$ nodes of the second type and

(3) $k^3 \sum_{t=1}^{k/4} \overline{\Psi}_{k-4t}(t+2)$ nodes of the third type.

By Corollary 1, we can thus conclude that there are at most $O(N/\log N)$ such nodes altogether. ∎

5. A Near-Optimal Layout

We are now prepared to describe a near-optimal, preliminary version of the optimal layout. In Section 6, we will show how to modify this layout in order to construct an optimal $O(N^2/\log^2 N)$-area layout for the N-node shuffle–exchange graph.

5a. Location of the Nodes

The near-optimal layout is constructed from a $\log N \times O(N/\log N)$ grid of nodes. Each column of the grid corresponds to a necklace of the shuffle–exchange graph. The nodes of each necklace are ordered from top to bottom so that the ith node is a left cyclic shift of the $(i - 1)$th node for each i and so that the distinguished node is placed in the bottom row. The necklaces are ordered from left to right so that the values of the distinguished nodes form an increasing sequence. (The *value* of a node is simply the numerical value of the associated k-bit binary string.) For example, we have constructed such a grid for the 32-node shuffle–exchange graph in Fig. 4. In the figure, we have represented each node in terms of the associated distinguished node. This representation readily illustrates the fact that the last bit of any node in the ith row corresponds to the ith bit of the associated distinguished node. Note that the necklaces $\langle 00000 \rangle$ and $\langle 11111 \rangle$ have not been included since they are degenerate.

5b. Insertion of the Edges

It is easily observed that the shuffle edges can be inserted in the grid with the addition of $O(N/\log N)$ vertical and 2 horizontal tracks. In the following, we will
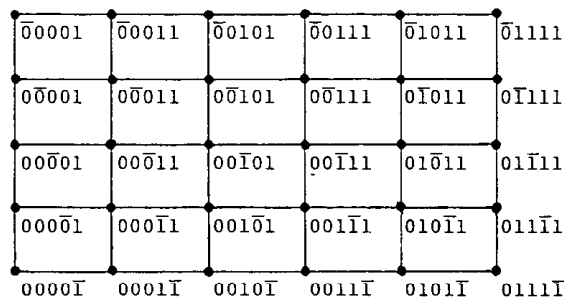
$\overline{0}0001$	$\overline{0}0011$	$\overline{0}0101$	$\overline{0}0111$	$\overline{0}1011$	$\overline{0}1111$
$0\overline{0}001$	$0\overline{0}011$	$0\overline{0}101$	$0\overline{0}111$	$0\overline{1}011$	$0\overline{1}111$
$00\overline{0}01$	$00\overline{0}11$	$00\overline{1}01$	$00\overline{1}11$	$01\overline{0}11$	$01\overline{1}11$
$000\overline{0}1$	$000\overline{1}1$	$001\overline{0}1$	$001\overline{1}1$	$010\overline{1}1$	$011\overline{1}1$
$0000\overline{1}$	$0001\overline{1}$	$0010\overline{1}$	$0011\overline{1}$	$0101\overline{1}$	$0111\overline{1}$

Fig. 4. The grid of nodes for the 32-node shuffle–exchange graph.

show that the exchange edges can be inserted with the addition of $O(N \log \log N/ \log N)$ vertical and horizontal tracks. Thus the total area of the layout is $O(N^2 \log \log^2 N/\log^2 N)$. This is only a factor of $O(\log \log^2 N)$ off from the lower bound of $O(N^2/\log^2 N)$.

The analysis is divided into two parts. First we show that only $O(N \log \log N/ \log N)$ exchange edges link nodes which are in *different* rows of the grid. Such edges can be inserted with the addition of at most $O(N \log \log N/\log N)$ vertical and horizontal tracks. We then conclude the analysis by showing that at most $O(N/\log N)$ horizontal tracks are needed to insert the exchange edges which link two nodes in the *same* row.

Consider an exchange edge which links two nodes that are in different rows of the grid. In particular, assume that the edge is incident to an even node in the ith row for some i. By definition, the even node can be represented as $w\bar{0}w'$, where $|w| = i - 1$ and $w0w'$ is the distinguished node of $\langle w0w' \rangle$. The exchange edge is also incident to the odd node $w\bar{1}w'$. By assumption, $w\bar{1}w'$ is not located in the ith row and thus $w1w'$ is *not* a distinguished node. Since $w0w'$ is a distinguished node, we know that the ith bit of $w0w'$ (the bit that was changed in order to produce $w\bar{1}w'$) must be in the primary or secondary block of zeros of $w0w'$. Otherwise, the primary and (if it exists) secondary blocks of zeros of $w1w'$ would be identical in location and size to the primary and secondary blocks of $w0w'$. This would imply that $w1w'$ is also distinguished, a contradiction. Thus $w\bar{0}w'$ must be a primary or secondary node. As was previously mentioned, we can assume that each necklace has at most $2 \log k = 2 \log \log N$ primary and $2 \log \log N$ secondary nodes. Thus at most $4 \log \log N$ nodes in each necklace are both even and incident to an exchange edge which links nodes in different rows. Since every exchange edge is incident to an even node and since there are $O(N/\log N)$ necklaces, we can concude that there are at most $O(N \log \log N/ \log N)$ exchange edges which link nodes in different rows.

We next show that those exchange edges which link two nodes that are in the same row can be inserted with the addition of at most $O(N/\log N)$ horizontal tracks. Once again, the analysis is divided into two parts. In the first part, we show that at most $O(N/\log N)$ exchange edges are contained in the first $\log k$ rows. Such edges can be trivially inserted with the addition of $O(N/\log N)$ horizontal tracks. In the second part, we show that only 2^{k-i} horizontal tracks are needed to insert the exchange edges in the ith row for any $i > \log k$. Since $\sum_{i=\log k + 1}^{k} 2^{k-i} \leqslant 2^k/k = N/\log N$, this will be sufficient to show that at most $O(N/\log N)$ additional horizontal tracks are necessary to insert the remaining exchange edges.

Consider a necklace which has t primary nodes for some $t \leqslant \log k$. By definition, the nodes in the first t rows of such a necklace are all even. Thus, such a necklace can have at most $r = \log k - t$ odd nodes in the first $\log k$ rows. By Lemma 3, we know that there are

$$\Psi_k(t)/k \sim (2^k/k)(e^{-k2^{-t-2}} - e^{-k2^{-t-1}})$$

such necklaces for $(\log k)/2 + \log \ln k \leqslant t \ll k$. By Lemma 4, we can assume that $t \geqslant$

$\log k - \log \ln k - 1$ and thus the total number of odd nodes occurring in the first $\log k$ rows is at most

$$\sim \sum_{t=\log k-\log\ln k-1}^{\log k} (\log k - t)(2^k/k)(e^{-k2^{-t-2}} - e^{-k2^{-t-1}})$$

$$= (2^k/k) \sum_{r=0}^{\log\ln k+1} r(e^{-k2^{r-2-\log k}} - e^{-k2^{r-1-\log k}})$$

$$= (2^k/k) \sum_{r=0}^{\log\ln k+1} r(e^{-2^{r-2}} - e^{-2^{r-1}})$$

$$\leqslant (2^k/k) \sum_{r=1}^{\log \ln k+1} e^{-2^{r-2}}$$

$$\leqslant (2^k/k) \sum_{r=1}^{\infty} e^{-2^{r-2}}$$

$$= O(N/\log N).$$

Since every exchange edge is incident to an odd node, the above bound implies that at most $O(N/\log N)$ exchange edges are contained in the first $\log k$ rows.

We next consider the number of horizontal tracks necessary to insert the exchange edges contained in the ith row for $i > \log k$. This number is identical to the maximum number of exchange edges that can overlap each other at a single point of the ith row. In Fig. 5, we illustrate the necessary conditions for two exchange edges to overlap in the ith row. All representations are in terms of distinguished nodes.

Note that the even end of an exchange edge is always to the left of the odd end. Also note that any node which occurs between $w\bar{0}w'$ and $w\bar{1}w'$ must be represented as $w\bar{0}w''$, where $w'' > w'$ or as $w\bar{1}w'''$, where $w''' < w'$. In either case, the exchange edge incident to the overlapped node extends beyond the exchange edge linking $w\bar{0}w'$ to $w\bar{1}w'$. Since there are at most $2^{k-i} - 1$ nodes between $w\bar{0}w'$ and $w\bar{1}w'$, these facts imply that at most 2^{k-i} exchange edges can overlap at any point of the ith row. This observation completes the argument that the near optimal layout requires only $O(N^2(\log\log N)^2/\log^2 N)$ area.
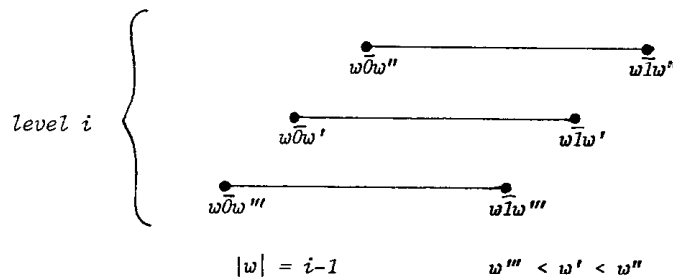

FIG. 5. Necessary conditions for exchange edges to overlap in the ith row.

6. An Optimal $O(N^2/\log^2 N)$-Area Layout

In this section, we will modify the layout described in Section 5 in order to produce an optimal $O(N^2/\log^2 N)$-area layout for the N-node shuffle–exchange graph. In particular, we will relocate the primary and secondary nodes of each necklace so that they are closer to and in the same row as the nodes to which they are linked via an exchange edge. Before going into the details of this relocation, however, it is necessary to introduce some additional terminology.

6a. More Definitions

In order to construct an optimal layout for the shuffle–exchange graph, we have found it necessary to break up each necklace into two or, possibly, three pieces. The *basic piece* of each necklace consists of all those nodes which are neither primary nor secondary. The *primary piece* of each necklace consists of the primary nodes while the *secondary piece* consists of the secondary nodes (if there are any). For example, the basic piece of $\langle 01011 \rangle$ is $\{0\bar{1}011, 010\bar{1}1, 0101\bar{1}\}$, the primary piece is $\{\bar{0}1011\}$, and the secondary piece is $\{010\bar{1}1\}$.

It is also necessary to extend the notion of a distinguished node to include pieces of necklaces. The *distinguished node of a basic piece* is the same as the distinguished node of the associated necklace. The *distinguished node of a primary piece* of a necklace is that node of the necklace which becomes distinguished when we ignore the primary block of zeros (i.e., when we temporarily replace the primary block of zeros in each node of the necklace with an equal-length block of ones). Similarly, the *distinguished node of a secondary piece* of a necklace is that node which becomes distinguished when we ignore the secondary block of zeros. For example, 010110111 is the distinguished node of the basic pice of $\langle 010110111 \rangle$, 011011101 is the distinguished node of the primary piece, and 011101011 is the distinguished node of the secondary piece. Note that the distinguished nodes of the primary and secondary pieces of every necklace are necessarily odd nodes and thus are contained in the basic piece of the necklace.

It is important to note that some necklaces (such as $\langle 01111 \rangle$) have a distinguished node but do not have a distinguished node for the primary or secondary piece of the necklace. Fortunately, arguments such as those used to prove Lemmas 5 and 6 can be used to show that at most $O(N/\log N)$ nodes are contained in such necklaces. Thus, we can assume henceforth that every piece of every necklace has an associated distinguished node.

6b. Location of the Nodes

As in Section 5, the layout is constructed from a $\log N \times O(N/\log N)$ grid of nodes. Each column of the grid corresponds to a piece of a necklace. The nodes of each piece are arranged within a column so that a node of the form $a_{k-1} \cdots \bar{a}_{k-i} \cdots a_0$ (where $a_{k-1} \cdots a_0$ is assumed to be the distinguished node of the associated piece) is placed in the ith row of the grid. Note that nodes in the basic piece of any necklace (these include all odd nodes) are in the same row as they were

in the near-optimal layout described in Section 5. The columns are ordered from left to right so that the values of the distinguished nodes of the associated pieces form a nondecreasing sequence. For example, we have constructed such a grid for $k = 5$ in Fig. 6.

Note that the necklaces $\langle 00001 \rangle$, $\langle 00011 \rangle$, $\langle 00111 \rangle$, and $\langle 01111 \rangle$ have not been included in Fig. 6 since their associated primary pieces do not have distinguished nodes.

6c. Insertion of the Edges

As each necklace is broken up into at most four *contiguous* pieces in the modified grid (the basic piece may have been broken up into two contiguous pieces), the shuffle edges can be inserted with the addition of at most $O(N/\log N)$ vertical and horizontal tracks. In what follows, we will show that at most $O(N/\log N)$ vertical and horizontal tracks are needed to insert all of the exchange edges as well. Thus the area of the layout will be $O(N^2/\log^2 N)$, which is optimal.

As before, we divide the analysis of the exchange edges into two parts. We first show that at most $O(N/\log N)$ exchange edges link nodes which are in different rows of the grid. Such edges can thus be trivially inserted with the addition of at most $O(N/\log N)$ vertical and horizontal tracks. We then show that those exchange edges which link two nodes in the same row can be inserted with the addition of only $O(N/\log N)$ horizontal tracks. The arguments will be very similar to those in Section 5b.

Consider an exchange edge which links two nodes which are in different rows of the grid. Since only primary and secondary nodes have been relocated, we can conclude from the arguments of Section 5b that the even node which is incident to the edge is either a primary or secondary node. In what follows, we will show that the even node is, in fact, a primary node.

Assume for the purposes of contradiction that the even node is a secondary node. Then this node can be represented as $w\bar{0}w'$ where $w0w'$ is the distinguished node of the secondary piece of $\langle w0w' \rangle$ and $|w| = i - 1$ for some i. By definition, $w\bar{0}w'$ is located in the ith row of the grid and is linked to $w\bar{1}w'$ via the exchange edge. Since
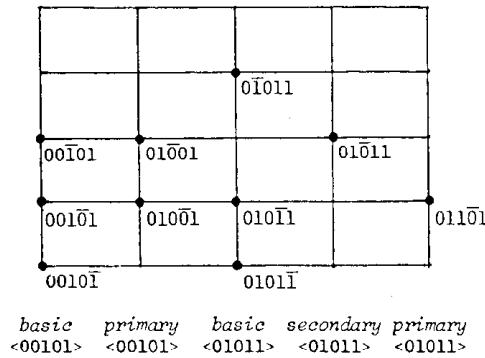

FIG. 6. Relocated nodes for the 32-node shuffle–exchange graph.

$w\bar{1}w'$ is odd, it is contained in the basic piece of $\langle w1w'\rangle$. By assumption, $w\bar{1}w'$ is not also in the ith row and thus $w1w'$ cannot be the distinguished node of $\langle w1w'\rangle$. Since the lengths of the two blocks of zeros in $w1w'$ created by switching the ith bit from 0 to 1 are less than the length of the primary block of zeros (in fact, the sum of their lengths is precisely one less than the length of the primary block), $w1w'$ will be the distinguished node of $\langle w1w'\rangle$ precisely when $w0w'$ is the node distinguished in $\langle w0w'\rangle$ by ignoring the secondary block of zeros. By definition, this is the case precisely when $w0w'$ is the distinguished node of the secondary piece of $\langle w0w'\rangle$. By assumption, $w0w'$ is the distinguished node of the secodary piece of $\langle w0w'\rangle$ and thus we can conclude that $w1w'$ is the distinguished node of $\langle w1w'\rangle$, a contradiction.

Next consider a *primary* node which is incident to an exchange edge linking two nodes in different rows of the grid. By the preceding arguments, this node must be of the form

$$w1\overbrace{0\cdots0}^{t_1}0\overbrace{0\cdots0}^{t_2}1w',$$

where $w10\cdots01w'$ is the distinguished node of the primary piece of $\langle w10\cdots01w'\rangle$ and either t_1 or t_2 is larger than or equal to the length of the longest block of zeros in $w11w'$. Otherwise,

$$w1\overbrace{0\cdots0}^{t_1}1\overbrace{0\cdots0}^{t_2}1w'$$

would (by definition) be the distinguished node of

$$\langle w1\overbrace{0\cdots0}^{t_1}0\overbrace{0\cdots0}^{t_2}1w'\rangle \qquad \text{and thus} \qquad w1\overbrace{0\cdots0}^{t_2}1\overbrace{0\cdots0}^{t_2}1w'$$

would be on the same row as

$$w1\overbrace{0\cdots0}^{t_1}\bar{0}\overbrace{0\cdots0}^{t_2}1w',$$

a contradiction. Each necklace contains at most $2r$ such primary nodes, where r is the difference between the lengths of the longest and second longest block of zeros in any string of the necklace. By Lemma 5, we can conclude that there are at most $O(N/\log N)$ such primary nodes in the entire shuffle–exchange graph. Thus, at most $O(N/\log N)$ exchange edges link nodes which are in different rows.

Using the analysis developed in Section 5b, it is not difficult to show that at most $O(N/\log N)$ horizontal tracks are needed to insert the exchange edges which link two nodes that are in the same row. In particular, there are still only $O(N/\log N)$ odd nodes in the top $\log k$ rows of the grid and thus at most $O(N/\log N)$ exchange edges are contained in the top $\log k$ rows. These can be trivially inserted with the addition of just $O(N/\log N)$ horizontal tracks.

Again following the methods of Section 5b, it is not difficult to show that two exchange edges overlap on the ith row only if the first i bits of the associated nodes are identical. Thus at most 2^{k-i} tracks are needed to insert all of the exchange edges in the ith row for all $i > \log k$. Summing, we can again conclude that at most $O(N/\log N)$ additional horizontal tracks are needed to insert the remaining exchange edges.

7. Layouts with Additional Edges

For some applications (such as the calculation of the discrete Fourier transform), it is useful to consider networks which have more than just shuffle and exchange edges. In particular, we might desire a layout for the shuffle–exchange graph which also includes shift, reverse and/or transpose edges. In what follows, we will show how to modify the optimal layout for the shuffle–exchange graph so that these additional edges can be inserted without increasing the total area by more than a constant factor.

7a. Shift Edges

Shift edges link the ith node to the $(i + 1)$th node for all *odd i*. When combined with the exchange edges, the resulting network will have links between the ith and the $(i + 1)$th nodes for all i. The inclusion of such edges facilitates the computation of discrete Fourier transforms at sequential intervals of a continuous signal. In such applications, the input data contained in the ith processor is shifted to the $(i + 1)$th processor for each i after each computation of a discrete Fourier transform. The graph consisting of shuffle, exchange and shift edges is known as the *shuffle–shift graph*.

Using the methods developed in Section 6, it is not difficult to show that the N-node shuffle–shift graph can be laid out using only $O(N^2/\log^2 N)$ area. As before, the necklaces are broken into two or three pieces and placed in a grid according to the value of the associated distinguished node. Thus the shuffle edges can be inserted as before using only $O(N/\log N)$ vertical and horizontal tracks.

For most odd nodes, adding a 1 to the value of the node changes only a relatively small number of bits at the end of the string. Thus it can be shown that at most $O(N/\log N)$ shift edges link nodes which are in different rows. These can be easily inserted using only $O(N/\log N)$ vertical and horizontal tracks. Of those edges which link nodes in the same row, at most $O(N/\log N)$ are contained in the first $\log k$ rows. For $i > \log k$, at most 2^{k-i} shift edges overlap at any point of the ith row. By introducing an extra vertical track for each necklace piece, it is possible to separate the layout of the shift edges on each level from that of the exchange edges. Thus both can be inserted simultaneously in the ith row using only $O(2^{k-i})$ total horizontal tracks. By the arguments of Section 6, this means that at most $O(N/\log N)$ additional horizontal tracks are needed to embed all of the remaining shift and exchange edges, thus completing the argument.

7b. *Reverse Edges*

Reverse edges link pairs of nodes that are associated with binary strings which are reverses of each other. For example, $a_{k-1} \cdots a_0$ is linked to $a_0 \cdots a_{k-1}$ via a reverse edge. Since the algorithm which computes discrete Fourier transforms on the shuffle–exchange network leaves the output for node $a_{k-1} \cdots a_0$ in node $a_0 \cdots a_{k-1}$, reverse edges provide a fast and convenient way of straightening out the solution. The graph consisting of shuffle, exchange, shift, and reverse edges will be referred to as the *shuffle–shift–reverse graph*.

Using the techniques developed in Section 6, it is also possible to show that the N-node shuffle–shift–reverse graph can be laid out in $O(N^2/\log^2 N)$ area. The basic idea is to modify the layout described in Section 7a so that

(1)   pieces of necklaces which are reverses of each other are paired together in the left-to-right ordering, and

(2)   pieces of necklaces are folded in half.

The first constraint insures that the maximal overlaps of the reverse edges in each row will be small while the second constraint insures that most reverse edges link nodes which are in the same row. Although it is not immediately obvious, it can be checked that these modifications do not substantially change the procedure for inserting the shuffle, shift and exchange edges which was described in Section 7a. Thus all of the edges can be inserted using at most $O(N/\log N)$ vertical and horizontal tracks.

7c. *Transpose Edges*

Transpose edges link the ith node to the $(N - 1 - i)$th node for each i. Viewed in terms of binary strings, transpose edges link each node to its complement. Although we do not know of any specific applications of transpose edges, they would be useful for problems that require frequent transposition of the data.

By further modifying the optimal layout for the shuffle–shift–reverse graph, it is possible to add transpose edges without increasing the total area by more than a constant factor. In particular, the layout should be modified so that

(1)   pieces of necklaces which are complements of each other are paired together in the left-to-right ordering, and

(2)   the distinguished node is selected on the basis of the location of the longest block of consecutive identical bits (be they zeros *or* ones).

The first constraint insures that the maximal overlaps of the transpose edges in each row are small while the second constraint insures that most transpose edges link nodes which are on the same row. Altough we do not present the details here, it is possible to show that such a layout can be constructed using only $O(N^2/\log^2 N)$ area, the least possible.

8. REMARKS

For some applications it is useful to consider optimizing measures other than area. For example, we might wish to minimize the number of wire crossings in the layout, the length of the longest wire in the layout, and/or the sum of the lengths of all the wires in the layout [10]. In [9], Leighton shows that any layout for the N-node shuffle–exchange graph must have $\Omega(N^2/\log^2 N)$ wire crossings. Since a layout with area A can have at most A wire crossings, the wire crossing lower bound is achieved by the layout described in this paper.

As a consequence of the wire crossing lower bound, it is also shown in [9] that any layout for the N-node shuffle–exchange graph must have an edge of length $\Omega(N/\log^2 N)$ and total wire length $\Omega(N^2/\log^2 N)$. The layout described in this paper clearly achieves the latter bound. It does not achieve the maximum wire length lower bound, however. In fact, we do not know of any layout for the N-node shuffle–exchange graph for which every wire has length $o(N/\log N)$. The layout described in this paper has wires of length $O(N/\log N)$.

The methods developed in this paper can be used to find several other optimal layouts for the shuffle–exchange graph. The key variant is the method by which a node is distinguished. In particular, the method must be impervious to small alterations in the necklace. (This is so that most exchange edges will link nodes which are in the same row of the grid.) Only by changing the values of a bit in a small segment of the necklace (such as the primary or secondary block of zeros) should we be able to globally change the distinguished node.

One such method of distinguishing a node is to select that node in the necklace which has the minimal value. Although the proof is fairly difficult, it can be shown that the layout for the N-node shuffle–exchange graph constructed in this manner has at most $O(N^2/\log^2 N)$ area.

As it previously was not known whether or not the N-node shuffle–exchange graph could be laid out in $O(N^2/\log^2 N)$ area, several researchers have tried to develop alternate networks which can compute discrete Fourier transforms in $O(N/\log N)$ steps and which can be laid out in $O(N^2/\log^2 N)$ area. The only other network discovered which has these properties is the *cube-connected-cycles graph* of Preparata and Vuillemin [17]. At this point, it is not clear which network best serves as the basis for practical parallel computation. Whereas the shuffle–exchange graph appears to be simpler to program than the cube-connected-cycles, the cube-connected cycles has somewhat simpler layouts. And although the shuffle–exchange graph appears to have smaller layouts for small values of N (see [13]), the cube-connected-cycles layouts are more regular and nicely structured.

Last, we would like to mention that the problem of finding a 3-dimensional layout for the shuffle–exchange graph with minimum volume remains unsolved. In fact, the shuffle–exchange graph is one of the few natural structures for which optimal 3-dimensional layouts are not known. (See [11, 14] for a discussion of general layout strategies in 2 and 3 dimensions.)

ACKNOWLEDGMENTS

REFERENCES

1. P. DIACONIS, R. L. GRAHAM, AND W. M. KANTOR, The mathematics of perfect shuffles, preprint, 1981.
2. L. J. GUIBAS AND A. M. ODLYZKO, Periods in strings, *J. Combin. Theory Ser. A* **30** (1981), 19–42.
3. L. J. GUIBAS AND A. M. ODLYZKO, String overlaps, pattern matching and nontransitive games, *J. Combin. Theory Ser. A* **30** (1981), 183–208.
4. D. HOEY AND C. E. LEISERSON, A layout for the shuffle–exchange network, *in* "Proceedings of the 1980 IEEE International Conference on Parallel Processing," August 1980.
5. D. KLEITMAN, F. T. LEIGHTON, M. LEPLEY, AND G. L. MILLER, New layouts for the shuffle–exchange graph, *in* "Proceedings of the 13th Annual ACM Symposium on Theory of Computing," pp. 278–292, May 1981.
6. T. LANG, Interconnection between processing and memory modules using the shuffle–exchange network, *IEEE Trans. Comput.* **C-25** (1976), 55–66.
7. C. E. LEISERSON, Area-efficient graph layouts for VLSI, *in* "Proceedings of the 21st Annual IEEE Symposium on Foundations of Computer Science," pp. 270–281, October 1980.
8. C. E. LEISERSON, "Area Efficient VLSI Computation," Ph.D. thesis Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Pa., November 1980.
9. 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 Mass., September 1981.
10. F. T. LEIGHTON, New lower bound techniques for VLSI, *Math. Systems Theory*, to appear.
11. F. T. LEIGHTON, A layout strategy for VLSI which is probably good *in* "Proceedings of the 14th Annual ACM Symposium on Theory of Computing," pp. 85–98, May 1982.
12. F. T. LEIGHTON, M. LEPLEY, AND G. L. MILLER, Layouts for the shuffle–exchange graph based on the complex plane diagram, *SIAM J. Algebraic and Discrete Meth.*, to appear.
13. F. T. LEIGHTON AND G. L. MILLER, Optimal layouts for small shuffle–exchange graphs, *in* "VLSI 81—Very Large Scale Integration" (J. P. Gray, Ed.), pp. 289–299, Academic Press, New York/ London, 1981.
14. F. T. LEIGHTON AND A. L. ROSENBERG, "Three-Dimensional Circuit Layouts," MIT-VLSI Technical Memo, No. 82–102.
15. D. S. PARKER, Notes on shuffle/exchange-type switching networks, *IEEE Trans. Comput.* **C-29** (3) (1980), 213–222.
16. F. P. PREPARATA, Optimal three-dimensional VLSI layouts, *Math. Systems Theory*, to appear.
17. F. P. PREPARATA AND J. E. VUILLEMIN, The cube-connected-cycles: A versatile network for parallel computation, *in* "Proceedings of the 20th Annual IEEE Symposium on the Foundations of Computer Science," pp. 140–147, October 1979.
18. H. S. STONE, Parallel processing with the perfect shuffle, *IEEE Trans. Comput.* **C-20** (2) (1971), 153–161.
19. J. T. SCHWARTZ, Ultracomputers, *ACM Trans. Programming Languages Systems* **2** (4) (1980), 484–521.
20. L. SNYDER, Overview of the CHiP computer, *in* "VSLI 81—Very Large Scale Integration" (J. Gray, Ed.), pp. 237–246, Academic Press, London/New York, 1981.

21. D. STEINBERG AND M. RODEH, A layout for the shuffle–exchange network with $O(N^2/\log^{3/2} N)$ area, *IEEE Trans. Comput.*, to appear.

22. C. D. THOMPSON, Area-time complexity for VLSI, *in* "Proceedings of the 11th Annual ACM Symposium on Theory of Computing," pp. 81–88, May 1979.

23. C. D. THOMPSON, "A Complexity Theory for VLSI," Ph.D. dissertation, Department of Computer Science, Carnegie-Mellon University, Pittsburg, Pa., 1980.