

A Layout for the Shuffle-Exchange Network

Dan Hoey
Charles E. Leiserson

Department of Computer Science
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

18 August 1980

Abstract

This paper describes a technique for producing a VLSI layout of the shuffle-exchange graph. It is based on the layout procedure in [2] which lays out a graph by bisecting the graph, recursively laying out the two halves, and then combining the two sublayouts. The area of the layout is related to the number of edges that must be cut to bisect the graph.

For the shuffle-exchange graph on n vertices, we present a bisection schema for which the above procedure yields an $O(n^2/\lg n)$ area layout when $n = 2^k$ and k is a power of two. The bisection involves a mapping from vertices of the graph to polynomials, and the polynomials are subsequently evaluated at complex roots of unity. Incidental to this construction is a result on the combinatorial problem of necklace enumeration.

This paper will appear in the Proceedings of the 1980 International Conference on Parallel Processing. Copyright © 1980, IEEE. Reproduction is permitted for all purposes of the U. S. Government.

This research was supported in part by the Defense Advanced Research Projects Agency (DoD), Arpa Order No. 3597, monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1551, the National Science Foundation under Grant MCS 78-236-76, and the Office of Naval Research under Contracts N00014-76-C-0370 and N00014-80-C-0236. Charles E. Leiserson receives support from a Fannie and John Hertz Foundation fellowship.

1. Introduction

The shuffle-exchange network has been shown to be an important communications structure for parallel processors. Stone [8] describes algorithms which use this structure to solve several problems, including the computation of the discrete Fourier transform and sorting bitonic sequences. The number of communications steps required by these algorithms is typically a polynomial in the logarithm of the number of nodes in the network, and the nodes themselves need only perform relatively simple operations.

VLSI designers often try to minimize the area used by a circuit subject to the requirements imposed by the fabrication technology on the minimum feature sizes of the components [5]. In [9] Thompson develops lower bounds on the growth of circuit area based on graph-theoretic properties of the communications structure. He shows in particular that any layout of the shuffle-exchange network on $n = 2^k$ vertices must use at least $\Omega(n^2/k^2)$ area. The arguments for Thompson's lower bounds are based on the *minimum bisection width* of a graph, which is the least number of edges that must be removed to separate the graph into two equal-sized subgraphs.

The concept of bisection width was extended by Lipton and Tarjan [3] to that of a *separator theorem* for a class of graphs closed under the subgraph relation. In essence, a separator theorem for a class provides upper bounds on the bisection widths of graphs in the class. Separator theorems allow the divide-and-conquer paradigm to be exploited in the design efficient algorithms for graph manipulation [4]. Recently, Leiserson [2] has used this approach to design area-efficient VLSI layouts.

In this paper a theorem similar to a separator theorem is proven for the shuffle-exchange graph on $n = 2^k$ vertices. We exhibit a *dissection* that shows how the shuffle-exchange graph may be bisected, how the resultant subgraphs may themselves be bisected, and so forth. We use this result to construct an $O(n^2/k)$ area layout for the case when k is a power of two, thereby improving Thompson's upper bound of $O(n^2/\sqrt{k})$. In our proof the vertices of the shuffle-exchange graph are mapped to a polynomial space, and then the polynomials are mapped to the complex plane. This construction also provides an asymptotic result on the combinatorial problem of necklace enumeration.

The next section formalizes the notions of bisection and dissection. Section 3 introduces the shuffle-exchange graph and describes its relationship to polynomials. In Section 4 we construct a bisection of the shuffle-exchange graph whose width is $O(n/k)$, and in Section 5 we extend this result to produce a dissection. In Section 6, the layout algorithm of [2] is applied to this dissection to produce an $O(n^2/k)$ area layout for the shuffle-exchange graph. Section 7 concludes by comparing this result with other work in the field.

2. Graph Dissection

In this section, we formalize concepts pertaining to the partitioning of a graph into smaller graphs by the removal of edges.

A *bisection* S of a graph $G = (V, E)$ into graphs $G' = (V', E')$ and $G'' = (V'', E'')$ is a disjoint partition of the vertices $V = V' \cup V''$ together with a disjoint partition of the edges $E = E' \cup E'' \cup E_S$ such that the cardinalities of V' and V'' differ by at most one. The cardinality of E_S is called the *width* of the bisection, and the edges in E_S are said to be *removed* by the bisection. The graphs G' and G'' are called the *halves* of the bisection.

Of course, any graph can be bisected by removing all its edges, but usually we are interested in removing as few edges as possible. The *minimum bisection width* of a graph is the smallest number of edges that must be removed to divide an n -vertex graph into a $\lceil n/2 \rceil$ -vertex graph and a $\lfloor n/2 \rfloor$ -vertex graph. Unfortunately, the problem of finding the minimum bisection width of an arbitrary graph is NP-complete [1].

It is sometimes the case that every graph in a class of graphs can be bisected by the same general mechanism. We define a *separator* for a class \mathcal{G} of graphs to be a family \mathcal{J} of bisections such that \mathcal{J} contains a bisection of every nontrivial graph G in \mathcal{G} . Interesting separators are those that exhibit the *closure property*. A separator \mathcal{J} for a class of graphs \mathcal{G} has this property if for any graph $G \in \mathcal{G}$, the halves G' and G'' that are produced by a bisection of G in \mathcal{J} are also in \mathcal{G} . Any separator with the closure property whose associated class contains a particular graph G is called a *dissection* of G .

A dissection \mathcal{J} of G may be thought of as a complete binary tree that has G at the root, the halves of G from some bisection in \mathcal{J} as its sons, and the halves of the halves as grandsons, and so forth to trivial graphs at the leaves. If G has n vertices, then the subgraphs at level j will have about $n/2^j$ vertices. Although there may be other graphs in the class \mathcal{G} associated with \mathcal{J} , at the very least \mathcal{G} must contain all of the graphs in the tree.

In [3] Lipton and Tarjan introduce *separator theorems* which use ideas similar to those presented here. In their work, however, the discussion is restricted to classes of graphs that are closed under the subgraph relation. (A class \mathcal{G} is closed under the subgraph relation if every subgraph G' of a graph $G \in \mathcal{G}$ is also an element of \mathcal{G} .) We have departed from their approach because the results of this paper rely on properties of the shuffle-exchange graph that do not hold for all of its subgraphs.

3. The Shuffle-Exchange Graph

The *shuffle-exchange graph* on n vertices is defined only when n is a power of two. Each of the $n = 2^k$ vertices can be identified with an element of the Cartesian product

$$\{0, 1\}^k = \{b_{k-1} b_{k-2} \cdots b_1 b_0 \mid b_j \in \{0, 1\}\}.$$

Each vertex $v \in \{0, 1\}^k$ is incident on an *exchange edge* $(v, \epsilon(v))$ and two *shuffle edges* $(v, \sigma(v))$ and $(v, \sigma^{-1}(v))$, where ϵ and σ are permutations defined by

$$\epsilon(b_{k-1} b_{k-2} \cdots b_1 b_0) = b_{k-1} b_{k-2} \cdots b_1 (1-b_0), \quad (1)$$

$$\sigma(b_{k-1} b_{k-2} \cdots b_1 b_0) = b_{k-2} b_{k-3} \cdots b_1 b_0 b_{k-1}. \quad (2)$$

In the literature the vertices are usually identified with integers from zero to $n-1$ represented in binary notation. The shuffle permutation σ is then the permutation applied to a deck of n cards by a perfect riffle shuffle, in which case $\sigma(m) \equiv 2m \pmod{n-1}$. The exchange permutation ϵ is the permutation that exchanges pairs of adjacent elements of the vertex set, so that $\epsilon(m) = m \pm 1$.

The shuffle-exchange graph is highly structured because of the shuffle permutation. From equation (2) we see that $\sigma(v)$ can be determined from v by rotating the indices of v to the left one position. The shuffle permutation partitions $\{0, 1\}^k$ into equivalence classes known as *necklaces* [7], where two vertices are equivalent whenever the indices of one are a cyclic permutation of the indices of the other. Since rotation by k positions yields the original vertex, the cardinality of a necklace cannot exceed k .

The properties that we shall use to dissect the shuffle-exchange graph are expressed conveniently in terms of the *characteristic polynomial*, which is defined for a vertex $v = b_{k-1} \cdots b_0 \in \{0, 1\}^k$ as

$$p_v(x) = \sum_{0 \leq j \leq k-1} b_j x^j. \quad (3)$$

It should be apparent that $p_v(2)$ is precisely the vector v considered as a binary number, as discussed above. The following lemma shows the relationship between the characteristic polynomial and the shuffle and exchange permutations.

Lemma 1: For all $v \in \{0, 1\}^k$,

$$p_{\epsilon(v)}(x) = p_v(x) \pm 1, \quad (4)$$

$$p_{\sigma(v)}(x) \equiv x p_v(x) \pmod{x^k - 1}, \quad (5)$$

where the congruence (5) is taken over the ring $\mathbb{Z}[x]$ of polynomials with integer coefficients.

Proof. From the defining equations (1) and (2),

$$p_v(x) - p_{\varepsilon(v)}(x) = b_0 x^0 - (1-b_0)x^0 = 2b_0 - 1,$$

$$x p_v(x) - p_{\sigma(v)}(x) = b_{k-1} x^k - b_{k-1} x^0 = b_{k-1} (x^k - 1).$$

The lemma follows from the fact that each b_j is either zero or one. \square

The cyclic structure of necklaces is exploited in Section 4 to bisect the shuffle-exchange graph. This is done in such a way that most of the necklaces in the graph are bisected. When the number of vertices in a necklace is even, it turns out that the half-necklaces also have a cyclic structure. An m -cycle is defined to be an ordered sequence $(v_0, v_1, \dots, v_{m-1})$ of m distinct vertices such that for $j = 1, \dots, m-1$,

$$p_{v_j}(x) \equiv x p_{v_{j-1}}(x) \pmod{x^m-1}. \quad (6)$$

The next lemma provides justification for calling such a sequence an m -cycle.

Lemma 2: Let (v_0, \dots, v_{m-1}) be an m -cycle. Any sequence $(v_j, \dots, v_{m-1}, v_0, \dots, v_{j-1})$ formed by cyclically permuting (v_0, \dots, v_{m-1}) is also an m -cycle. If d is a divisor of m , then the subsequence (v_0, \dots, v_{d-1}) is a d -cycle.

Proof. This lemma can be proved by manipulating the congruence (6) in the definition of an m -cycle. The congruence can be iterated to yield

$$p_{v_{m-1}}(x) \equiv x^{m-1} p_{v_0}(x) \pmod{x^m-1},$$

and since $x^m \equiv 1 \pmod{x^m-1}$, it follows that

$$x p_{v_{m-1}}(x) \equiv p_{v_0}(x) \pmod{x^m-1}.$$

Thus (6) holds between the first and last vertices as well as between adjacent vertices, implying that the choice of a first vertex is immaterial. To prove the second part of the lemma, observe that congruence (6) modulo x^m-1 must also hold modulo its divisor x^d-1 . \square

Congruence (5) shows that a necklace of k vertices is a k -cycle. Lemma 2 establishes that when k is even, the necklace can be bisected to yield two $k/2$ -cycles.

4. Bisecting the Shuffle-Exchange Graph

The concepts developed in Section 3 are applied in this section to construct a bisection of the shuffle-exchange graph on $n = 2^k$ vertices. The construction is obtained by evaluating the characteristic polynomials of the vertices at a complex k th root of unity, inducing a mapping from $\{0, 1\}^k$ to the complex plane. The complex plane is then divided to induce a bisection of the shuffle-exchange graph. A corollary of this construction is an asymptotic result on the number of necklaces.

Let $\omega = e^{2\pi i/k}$ be the principal primitive complex k th root of unity, and consider the mapping $v \mapsto p_v(\omega)$ from $\{0, 1\}^k$ to the complex plane. Figure 1 graphs the values of $p_v(\omega)$ for $k = 5$. The vertices are labeled with $p_v(2)$. The solid lines forming pentagons concentric about the origin represent shuffle edges, and the horizontal dotted arcs represent exchange edges.




Figure 1: The shuffle-exchange graph on $32 = 2^5$ vertices mapped to the complex plane by $v \mapsto p_v(\omega)$. Vertices are labelled with $p_v(2)$. Dotted lines represent exchange edges, and solid lines represent shuffle edges.

Let us examine this figure in relation to Lemma 1. The occurrence of regular k -gons of shuffle edges can be explained by congruence (5). Since ω is a root of $x^k - 1$, this congruence becomes the equality $p_{\sigma(v)}(\omega) = \omega p_v(\omega)$. Thus $p_{\sigma(v)}(\omega)$ is the point obtained from $p_v(\omega)$ by a counterclockwise rotation of $2\pi/k$ radians about the origin. The vertices in a necklace are mapped to k points equally spaced on a circle about the origin, unless the entire necklace is mapped to the origin. The fact that exchange edges are horizontal can be explained by equation (4) in Lemma 1. If vertices v and $\epsilon(v)$ are incident on an exchange edge, then they are mapped to complex numbers that have the same imaginary part and differ by one in the real part.

The bisection of the shuffle-exchange graph will be achieved by partitioning the vertices based on the imaginary part of $p_v(\omega)$, with tie-breaking when $p_v(\omega)$ is real. All edges that cross the real line will be removed, and it will be shown that there are at most $O(n/k)$ of these. This bound is easily shown for edges whose incident vertices are not involved in the tie-breaking. Since there are n vertices in the shuffle-exchange graph, there are at most n/k regular k -gons of shuffle edges, and each of these k -gons crosses the real line twice. Since exchange edges are horizontal, they never cross the real line.

In order to define the bisection formally, we first partition the nonzero complex numbers as $\mathbb{C}^+ \cup \mathbb{C}^-$ where

$$\begin{aligned}\mathbb{C}^+ &= \{z \in \mathbb{C} \mid \text{Im}(z) > 0\} \cup \{x \in \mathbb{R} \mid x > 0\}, \\ \mathbb{C}^- &= \{z \in \mathbb{C} \mid \text{Im}(z) < 0\} \cup \{x \in \mathbb{R} \mid x < 0\}.\end{aligned}$$

The halves G' and G'' are defined by the regions to which vertices of the shuffle-exchange graph are mapped. The vertices for which $p_v(\omega) \in \mathbb{C}^+$ are assigned to V' and those for which $p_v(\omega) \in \mathbb{C}^-$ are assigned to V'' . The remaining vertices, those for which $p_v(\omega) = 0$, are distributed arbitrarily but equally between V' and V'' . Three types of edges are placed in E_S

1. Exchange edges whose incident vertices are mapped to real numbers.
2. Shuffle edges whose incident vertices are mapped to the origin.
3. Shuffle edges between vertices v and v' such that $p_v(\omega) \in \mathbb{C}^+$ and $p_{v'}(\omega) \in \mathbb{C}^-$.

It can be seen by inspection that E_S is a superset of the set of edges that connect V' to V'' . Edges not in E_S are allocated to E' or E'' according as their incident vertices are in V' or V'' .

To see that $|V'| = |V''|$, consider for any vertex v the vertex $\mathcal{C}(v)$ obtained by complementing every index in the vector v . This relationship can be restated in terms of characteristic polynomials as

$$p_{\mathcal{C}(v)}(x) = (x^{k-1} + x^{k-2} + \dots + 1) - p_v(x).$$

Because the sum of all k th roots of unity is zero, it follows that $p_v(\omega) = -p_{\mathcal{C}(v)}(\omega)$. Therefore, the correspondence $v \leftrightarrow \mathcal{C}(v)$ is a one-to-one correspondence between the vertices mapped to \mathbb{C}^+ and those mapped to \mathbb{C}^- . This proves that this partition is a bisection as was claimed. The cardinality of E_S is the width of the bisection and is bounded by the following theorem.

Theorem 3: For any positive integer k , there is a bisection S of the shuffle-exchange graph on $n = 2^k$ vertices such that the width of S is at most $6(n/k)$.

Proof. Let S be the bisection described above, and consider the three types of edges that compose E_S . We will bound each of the three types by the quantity $2(n/k)$.

Each of the type 3 edges is a shuffle edge incident on vertices mapped to nonzero complex numbers, and each such vertex belongs to a necklace of exactly k vertices which are mapped to nonzero numbers. Since the total number of vertices in the shuffle-exchange graph is n , there can be at most n/k such necklaces. The

shuffle edges in each of these necklaces form a regular k -gon centered at the origin, and thus only two of these edges can cross the real line, in the sense of having one incident vertex mapped to \mathbb{C}^+ and the other to \mathbb{C}^- . Thus there can be at most $2(n/k)$ type 3 edges.

The same argument can be used to bound the number of type 1 edges. There are at most $2(n/k)$ vertices mapped to nonzero real numbers. Since every exchange edge whose incident vertices are mapped to real numbers has at least one of these vertices mapped to a nonzero real number, there can be no more than $2(n/k)$ type 1 edges.

Finally, the number of type 2 edges can be bounded by the number of type 1 edges by observing that for each shuffle edge $(v, \sigma(v))$ whose incident vertices are mapped to the origin, the exchange edge $(\sigma(v), \epsilon(\sigma(v)))$ is a type 1 edge. \square

We now pause to examine an interesting by-product of these counting arguments, a result on the combinatorial problem of necklace enumeration. A necklace is a string of k pearls, where each pearl may be one of c colors. Two necklaces are considered equivalent if one can be rotated to form the other, but not if they are only reflections. It is well-known [7] that the number of necklaces of k pearls in c colors is

$$(1/k) \sum_{d|k} c^{k/d} \phi(d). \quad (7)$$

In this formula $\phi(d)$ is Euler's totient function, the number of positive integers not exceeding d that are relatively prime to d . Although it appears that the term for $d=1$ in (7) might dominate the summation, it is not apparent that the contribution of the other terms is insignificant. However, the following corollary to Theorem 3 shows that this term is asymptotically dominant.

Corollary: The number of necklaces of $\{0, 1, \dots, c-1\}^k$ lies between c^k/k and $((c+1)/(c-1))(c^k/k)$.

Proof. The definitions of the σ and ϵ permutations may be extended to $\{0, 1, \dots, c-1\}^k$ as follows.

$$\epsilon(b_{k-1} b_{k-2} \dots b_1 b_0) = b_{k-1} b_{k-2} \dots b_1 (b_0+1 \bmod c),$$

$$\sigma(b_{k-1} b_{k-2} \dots b_1 b_0) = b_{k-2} \dots b_1 b_0 b_{k-1}.$$

The characteristic polynomial is defined as before (notice that now $p_v(c)$ is the vector v considered as a number expressed in base c notation), and the argument of Theorem 3 can be adapted to show that the function $v \mapsto p_v(\omega)$ maps at most $2c^k/(c-1)k$ elements of $\{0, 1, \dots, c-1\}^k$ to zero and that the remainder lie in necklaces of k elements. \square

5. Dissecting the Shuffle-Exchange Graph

In the previous section, we presented a bisection of the shuffle-exchange graph on $n = 2^k$ vertices. In this section we will show that when k is even, the structure of the halves is similar to the structure of the original shuffle-exchange graph. This similarity is captured in the notion of an *m-cyclic subgraph* of the shuffle-exchange graph, and it is shown that the halves are $k/2$ -cyclic subgraphs. The bisection from Theorem 3 can be modified to bisect *m-cyclic subgraphs* when m is even. Thus when k is a power of two, this approach can be used iteratively to construct a complete dissection of the shuffle-exchange graph.

An *m-cyclic subgraph* is a subgraph of the shuffle-exchange graph whose vertices are partitioned into disjoint *m-cycles*. Vertices not appearing in these *m-cycles* are also allowed, but such vertices must be isolated, not incident on any edge in the subgraph. If a shuffle edge $(v, \sigma(v))$ appears as an edge of the *m-cyclic subgraph*, it must occur between adjacent vertices of one of the *m-cycles*, and the exchange edge $(\sigma(v), \epsilon(\sigma(v)))$ must be an edge of the *m-cyclic subgraph* as well.

The reader should be warned that *m-cyclic subgraphs* are nothing more than a vehicle for extending the bisection of the shuffle-exchange graph to a dissection. The definition has been carefully crafted so that the proof of Theorem 3 will apply to them and so that their separator exhibits the closure property.

Lemma 4: When k is even, the halves G' and G'' produced by the bisection from Theorem 3 are $k/2$ -cyclic subgraphs.

Proof. Without loss of generality, we show this for G' only. The vertices that are mapped to zero by $v \mapsto p_v(\omega)$ have no incident edges (are isolated), but every other vertex of G' occurs in some sequence $(v_0, \dots, v_{k/2-1})$ that arose from cutting a necklace of k vertices in half. Since any necklace of k vertices is a k -cycle, and $k/2$ divides k , Lemma 2 ensures that this sequence is a $k/2$ -cycle. Thus we have demonstrated the first requirement for G' to be an $k/2$ -cyclic subgraph: every vertex not in an *m-cycle* is isolated.

We must now show that if a shuffle edge $(v, \sigma(v))$ appears as an edge in G' , then it occurs between adjacent vertices of one of the *m-cycles*, and furthermore, that then the exchange edge $(\sigma(v), \epsilon(\sigma(v)))$ is also in G' . It is clear that the first condition is satisfied. The second condition can be demonstrated by observing that both v and $\sigma(v)$ are mapped to \mathbb{C}^+ . Since the point $p_{\sigma(v)}(\omega)$ can be obtained from $p_v(\omega)$ by a counterclockwise rotation of $2\pi/k < \pi$ radians about the origin, it is impossible for $\sigma(v)$ to be mapped to the real line. The set of removed edges E_S contains only those exchange edges whose incident vertices are mapped to real points, which means that $(\sigma(v), \epsilon(\sigma(v)))$ must be in E' . \square

When m is even, the bisection from Theorem 3 can be generalized to a bisection of an arbitrary *m-cyclic subgraph*. Let $\omega_m = e^{2\pi i/m}$ and consider the function $v \mapsto p_v(\omega_m)$. Since ω_m is a root of $x^m - 1$, the congruence (6) between adjacent vertices of *m-cycles* becomes the equality $p_{v_j}(\omega_m) = \omega_m p_{v_{j-1}}(\omega_m)$. This

means that if any vertex of an m -cycle is mapped to a nonzero complex number, all the m vertices of the m -cycle are mapped to distinct points evenly spaced on a circle about the origin. Equation (5) applies as before to show that vertices connected by an exchange edge are mapped to complex numbers which differ by one.

Let G be an arbitrary m -cyclic subgraph of a shuffle-exchange graph on $n = 2^k$ vertices, and suppose that m is even. In order to construct a bisection of G , the vertices of the m -cycles of G are assigned to V' or V'' according as they are mapped by $v \mapsto p_v(\omega_m)$ to \mathbb{C}^+ or \mathbb{C}^- . The remaining vertices of G are those vertices that are mapped to the origin and those that are isolated. These may be divided arbitrarily but equally between V' and V'' . As with the bisection from Theorem 3, E_S consists of three types of edges.

1. Exchange edges whose incident vertices are mapped to real numbers.
2. Shuffle edges whose incident vertices are mapped to the origin.
3. Shuffle edges between vertices v and v' such that $p_v(\omega_m) \in \mathbb{C}^+$ and $p_{v'}(\omega_m) \in \mathbb{C}^-$.

The remaining edges are assigned to E' or E'' depending on whether their incident vertices are in V' or V'' .

Unlike before, however, the correspondence $v \leftrightarrow \mathbb{C}(v)$ cannot be used to show that $|V'| = |V''|$, since v may be a vertex of G when $\mathbb{C}(v)$ is not. But because m is even, the equality $p_{v_j}(\omega_m) = -p_{v_{j+m/2}}(\omega_m)$ holds for vertices v_j and $v_{j+m/2}$ in the same m -cycle, and the correspondence $v \leftrightarrow v_{j+m/2}$ suffices to show that this partition is a bisection. The following lemma provides a bound for the width of the bisection.

Lemma 5: Let m be even, and let G be an m -cyclic subgraph on t vertices. There is a bisection S that bisects G into $m/2$ -cyclic subgraphs and has width at most $6t/m$.

Proof. Let S be the bisection just described. Its width can be bounded by showing that there are at most $2t/m$ of each of the three types of edges in E_S . This bound holds for type 3 edges because there can be at most t/m disjoint m -cycles in G and no more than two type 3 edges per m -cycle. Since each type 1 edge has at least one incident vertex mapped to a nonzero real number, and there are at most two such vertices per m -cycle, the bound holds for these edges. Finally, for any type 2 edge $(v, \sigma(v))$, the edge $(\sigma(v), \epsilon(\sigma(v)))$ is a type 1 edge because G is an m -cyclic subgraph. Thus there can be no more type 2 edges than type 1 edges, and the bound on the width of the bisection is proved. It should be remarked here that the definition of m -cyclic subgraphs was specifically constructed in order to establish this correspondence between type 1 and type 2 edges.

To prove that the halves of the bisection are $m/2$ -cyclic subgraphs, observe that the bisection S isolates those vertices that are in m -cycles mapped to the origin, and splits the other m -cycles into pairs of $m/2$ -cycles. Since shuffle edges appear only between adjacent vertices of m -cycles, this adjacency is preserved in the $m/2$ -cycles. The only exchange edges removed by the bisection are those whose incident vertices are mapped to real numbers, and hence the argument of Lemma 4 can be used to show that if $(v, \sigma(v))$ is in one of the halves, then $(\sigma(v), \epsilon(\sigma(v)))$ is also in the half. \square

We are now ready to combine this bisection with the bisection from Theorem 3 into a dissection of the shuffle-exchange graph on $n = 2^k$ vertices for the case when k is a power of two. Recall from Section 2 that to dissect this graph, we need to find a class of subgraphs that has a separator with the closure property. The next theorem provides such a class.

Theorem 6: If k is a power of two, then there is a dissection \mathcal{F}_n of the shuffle-exchange graph on $n = 2^k$ vertices such that any bisection in \mathcal{F}_n which bisects an m -vertex graph has width at most

$$f_n(m) = \begin{cases} 6n/k & \text{if } m > n/k, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Proof. Let \mathcal{G}_n be the class of subgraphs consisting of *i*) the shuffle-exchange graph itself, *ii*) its $k/2^j$ -cyclic subgraphs that have $n/2^j$ vertices, for $j = 1, \dots, (\lg k) - 1$, and *iii*) its subgraphs that have no edges. Correspondingly, the separator \mathcal{F}_n consists of *i*) the bisection of the shuffle-exchange graph from Theorem 3, *ii*) the bisections of its $k/2^j$ -cyclic subgraphs from Lemma 5, and *iii*) arbitrary bisections of the totally disconnected subgraphs. To see that the closure property holds for \mathcal{F}_n , we first observe that the halves of the shuffle-exchange graph are $k/2$ -cyclic subgraphs with $n/2$ vertices. For $j = 1, \dots, (\lg k) - 2$, the halves of the $k/2^j$ -cyclic subgraphs with $n/2^j$ vertices are $k/2^{j+1}$ -cyclic subgraphs with $n/2^{j+1}$ vertices. When $j = (\lg k) - 1$ the bisection from Lemma 5 uses the mapping $v \mapsto p_v(\omega_2)$ to bisect 2-cyclic subgraphs. Since $\omega_2 = -1$, all vertices are mapped to real numbers, and thus the halves consist entirely of isolated vertices.

The bisection of the shuffle-exchange graph from Theorem 3 has width $6(n/k)$. For $j = 1, \dots, (\lg k) - 1$, the bisection from Lemma 5 bisects a $k/2^j$ -cyclic subgraph of $n/2^j$ vertices with width $6(n/2^j)/(k/2^j) = 6(n/k)$. The totally disconnected graphs can be bisected with zero width. \square

6. Laying Out the Shuffle-Exchange Graph

Given a dissection of an arbitrary graph, the divide-and-conquer technique of [2] can produce a VLSI layout whose area is related to the bisection widths of the graphs in the dissection. The VLSI model used is that of [9], and its important attributes are that wires have a minimum width and that only a constant number may cross at a point. In this section the results of Section 5 are applied to produce an $O(n^2/\lg n)$ area VLSI layout for an n -vertex shuffle-exchange network.

The technique of [2] constructs a layout for a general graph G by first bisecting G and laying out the halves recursively. The halves are then placed side-by-side, and the edges that were removed to bisect G are routed between the halves. The layout area can therefore be described as a recurrence in the area of the halves and the area required to route the edges removed by the bisection. This latter quantity is a function of the bisection widths in the dissection of G because the length and width of the layout increase by a constant

amount for each edge routed.

The particulars of how the area recurrence arises from this construction are described more fully in [2]. Some solutions to the recurrence are also given in that paper, but the bisection width bound $f_n(m)$ from equation (8) fails to satisfy certain conditions that are assumed for those solutions. Therefore, we give the area recurrence from [2] without further justification, but present its solution in detail.

Let $A_n(m)$ be the area of the layout of an m -vertex graph in the dissection of Theorem 6 (thus $A_n(n)$ is the area of the original shuffle-exchange graph). We express $A_n(m)$ in terms of $f_n(m)$ from equation (8). For the initial condition of the area recurrence, $A_n(1)$ is a constant, and for $1 < m \leq n$,

$$A_n(m) = [\sqrt{2 A_n(m/2)} + f_n(m)]^2. \quad (9)$$

The recurrence can be solved by taking the square root of both sides and then substituting $L_n(m)$ for $\sqrt{A_n(m)}$. For $1 < m \leq n$ this yields

$$L_n(m) = \sqrt{2} L_n(m/2) + f_n(m).$$

Iterating this recurrence and recalling that $n = 2^k$, we have

$$\begin{aligned} L_n(n) &= f_n(n) + \sqrt{2} f_n(n/2) + 2 f_n(n/4) + \dots \\ &\quad + \sqrt{2}^{k-1} f_n(2) + \sqrt{2}^k f_n(1) + \sqrt{n} L_n(1) \\ &\leq (6n/k) [1 + \sqrt{2} + \dots + \sqrt{2}^{\lg k}] + \sqrt{n} L_n(1) \\ &= (6n/k) [\sqrt{2}^{(\lg k)+1} - 1] / (\sqrt{2} - 1) + \sqrt{n} L_n(1) \\ &= O((n/k) \sqrt{k}) \\ &= O(n/\sqrt{k}). \end{aligned} \quad (10)$$

The reason the sum of the powers of $\sqrt{2}$ goes only as far as $\lg k$ in line (10) is that $f_n(m)$ is zero after this point. Since $A_n(n)$ is the square of $L_n(n)$, the area of the layout is $O(n^2/k)$.

This technique has been used in Figure 2 to lay out a shuffle-exchange network on 256 vertices. Only one fourth of the layout is shown, and the dissection that was used differs slightly from the one in Section 5. Instead of removing exchange edges, the arbitrary divisions among vertices mapped to zero are chosen so that $\epsilon(v)$ is in the same component as v , and the two are placed together.

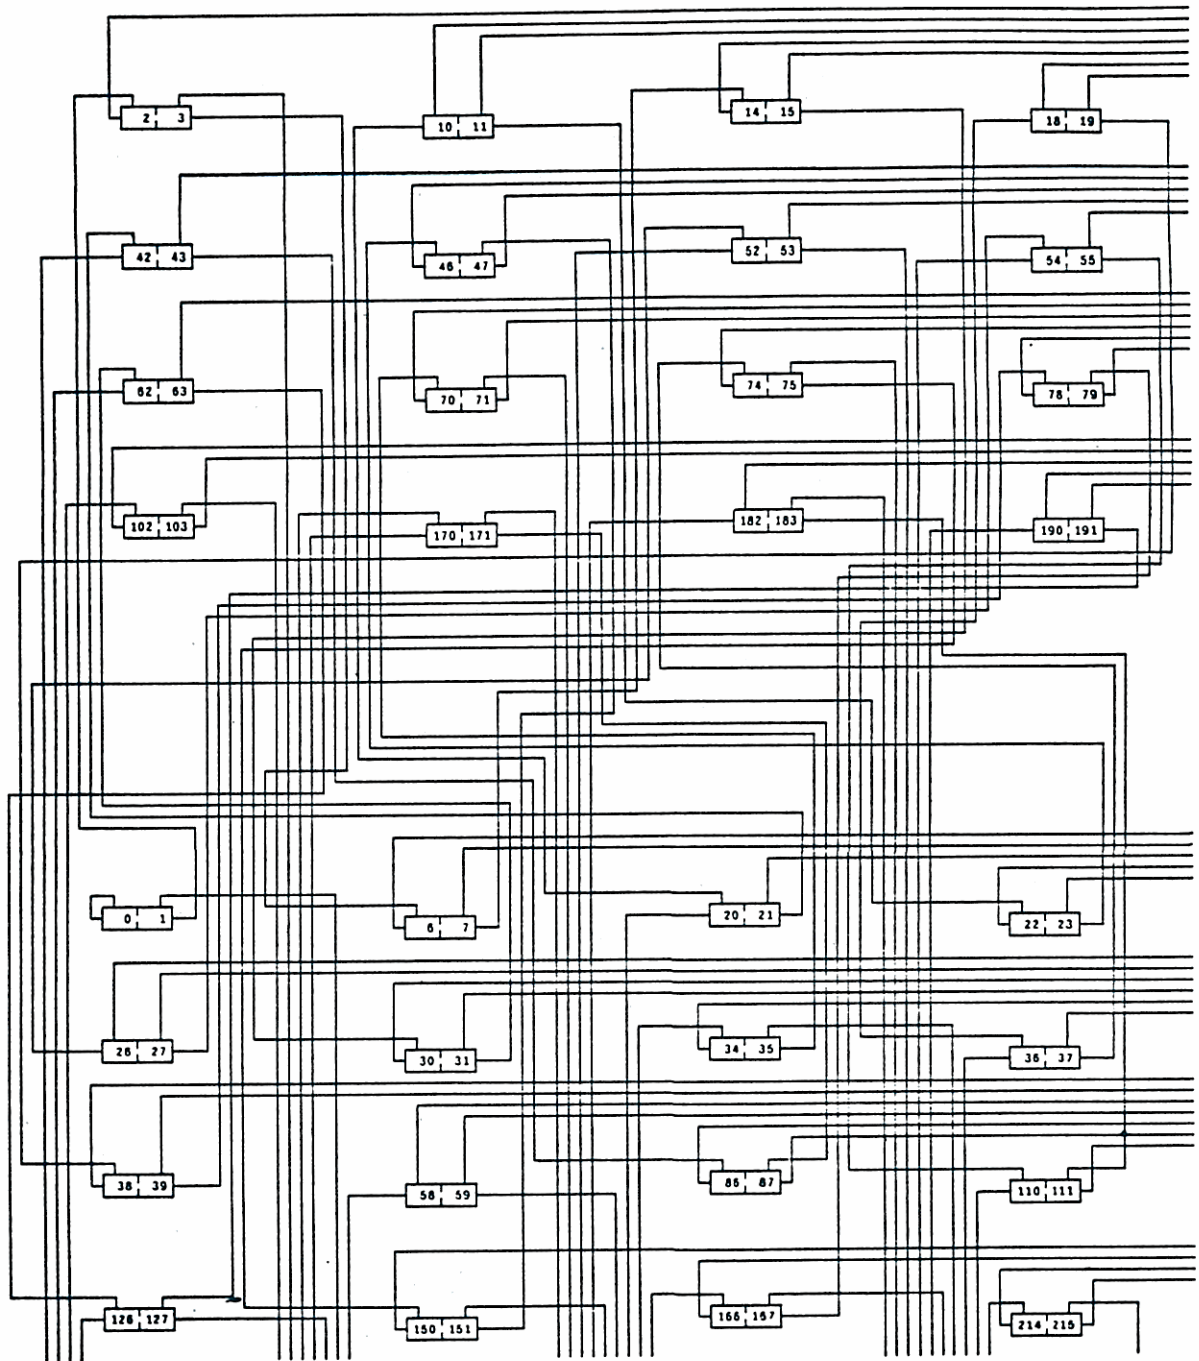


Figure 2: One fourth of a shuffle-exchange network.

7. Conclusion

We have developed an extraordinary amount of machinery in order to construct an $O(n^2/k)$ area layout for the shuffle-exchange graph on $n = 2^k$ vertices, and indeed, we have only been able to show this upper bound for the case when k is a power of two. It may be that this bound holds when k is not a power of two, but we have not been able to prove this. For the time being, the best general upper bound seems to be Thompson's $O(n^2/\sqrt{k})$ bound.

In any event, a gap remains between either of these upper bounds and the best known lower bound of $\Omega(n^2/k^2)$ which is also given by Thompson. This lower bound is proved in [9] by showing that the minimum bisection width of the shuffle-exchange graph must be $\Omega(n/k)$ and that the area of any graph layout must be at least the square of the minimum bisection width of the graph. Theorem 3 shows that this $\Omega(n/k)$ lower bound for bisection of the shuffle-exchange graph can be achieved, even though the dissection based on this bisection does not achieve the $\Omega(n^2/k^2)$ lower bound for layout area. This is because the bisection width $f_n(m)$ does not immediately decrease as m decreases from n . It may be that an improved lower bound for the layout area will be based on the notion of a minimum dissection, where the width of every bisection in any dissection can be bounded from below.


On the other hand, it may be that an $O(n^2/k^2)$ area layout does exist for the shuffle-exchange graph, as does one for the *cube-connected-cycles* (CCC) network of Preparata and Vuillemin [6]. The CCC is the graph that arises from a boolean hypercube of d dimensions when each vertex is replaced by a cycle of d vertices. Many of the problems that can be solved quickly using the shuffle-exchange interconnection can also be solved quickly using the CCC. But despite the fact that a smaller layout is known for the CCC, descriptions of algorithms for the CCC tend to be more complicated. The discovery of an $O(n^2/k^2)$ area layout for the shuffle-exchange graph would therefore favor the shuffle-exchange graph as the network of choice and would allow the many algorithms already designed for this network to be applied directly in optimal VLSI implementations. But until such a layout is found—if ever one is found—the CCC will continue to have the edge.

In conclusion, we believe that characteristic polynomials provide a useful way of viewing the shuffle-exchange network, and we believe that this approach goes beyond the particular technical results presented here. Characteristic polynomials unveil properties of the shuffle-exchange graph that are obscured by the classical approach of relating the vertices to integers. We hope that the mechanisms we have developed to relate the topology of a particular graph to the algebra of polynomials will be exploited further.

References

- [1] M. R. Garey, D. S. Johnson, and L. Stockmeyer, "Some simplified polynomial complete problems," *6th Annual Symposium on Theory of Computing*, ACM, (April, 1974), pp. 47-63.
- [2] C. E. Leiserson, "Area-efficient graph layouts (for VLSI)," *21st Annual Symposium on Foundations of Computer Science*, IEEE Computer Society, (October, 1980).
- [3] R. J. Lipton and R. E. Tarjan, "A separator theorem for planar graphs," *A Conference on Theoretical Computer Science*, University of Waterloo, (August, 1977).
- [4] R. J. Lipton and R. E. Tarjan, "Applications of a planar separator theorem," *18th Annual Symposium on Foundations of Computer Science*, IEEE Computer Society, (October, 1977), pp.162-170.
- [5] C. A. Mead and L. A. Conway. *Introduction to VLSI Systems*, Addison-Wesley, (1980).
- [6] F. P. Preparata and J. Vuillemin. *The cube-connected-cycles: a versatile network for parallel computation*, Technical Report 356, Institut de Recherche d'Informatique et d'Automatique, (June, 1979).
- [7] J. Riordan, *An Introduction to Combinatorial Analysis*, John Wiley & Sons, Inc., (1958).
- [8] H. S. Stone, "Parallel processing with the perfect shuffle," *IEEE Transactions on Computers*, C-20, 2, (February, 1971), pp.153-161.
- [9] C. D. Thompson, *A Complexity Theory for VLSI*, Ph.D. Thesis, Carnegie-Mellon University Computer Science Department, (1980).

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER CMU-CS-80-139	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) A LAYOUT FOR THE SHUFFLE-EXCHANGE NETWORK		5. TYPE OF REPORT & PERIOD COVERED Interim
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) DAN HOEY CHARLES E. LEISERSON		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0370
		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
9. PERFORMING ORGANIZATION NAME AND ADDRESS Carnegie-Mellon University Computer Science Department Pittsburgh, PA 15213		12. REPORT DATE AUGUST 1980
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Arlington, VA 22217		13. NUMBER OF PAGES 16
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		






Figure 3: A layout of the shuffle-exchange graph on 64 vertices. Vertical loops represent shuffle edges; horizontal lines represent exchange edges.

Addendum

The construction described in Section 6 requires $O(n^2 / \lg n)$ area to lay out the shuffle-exchange graph on $n = 2^k$ vertices when k is a power of two, but this technique cannot be used for other values of k . Here we provide a construction of an $O(n^2 / \lg n)$ area layout for all values of k . This improved construction does not use the results of Sections 5 and 6, but follows directly from the mapping $v \mapsto p_v(\omega)$ described in Section 4.

The idea behind the construction is simple. Figure 3 shows a layout of the shuffle-exchange graph on 64 vertices. Each exchange edge is allocated a constant-width track parallel to the x -axis. Each necklace is allocated two tracks parallel to the y -axis. The shuffle edges in the necklace run up one track and down the other. Since there are $O(n / \lg n)$ necklaces and $O(n)$ exchange edges, the area of the layout is $O(n^2 / \lg n)$.

For this construction, the assignment of x -coordinates to necklaces is arbitrary, but the y -coordinates of the vertices must satisfy two constraints. First, y -coordinates of vertices connected by an exchange edge must be equal because exchange edges run parallel to the x -axis. Second, the y -coordinates of the vertices $v, \sigma(v), \sigma(\sigma(v)), \dots, \sigma^{-1}(v)$ in a given necklace must form a *bitonic sequence* [8], that is, the coordinates must consist of an increasing sequence and a decreasing sequence. This condition guarantees that two tracks suffice to route the shuffle edges.

In order to assign y -coordinates in this way, recall the mapping $v \mapsto p_v(\omega)$ from Section 3. First, vertices connected by exchange edges are mapped to complex numbers with the same imaginary component. Second, the vertices in each necklace, if not all mapped to the origin, are mapped to complex numbers in order on a circle, which implies that the imaginary parts of these complex numbers form a bitonic sequence. Therefore, an assignment of y -coordinates in the same order as $\text{Im}(p_v(\omega))$, breaking ties arbitrarily, is consistent with the two requirements listed above.

Vertices mapped to the origin must also be assigned y -coordinates in a bitonic sequence. To do this, notice that ties can be broken arbitrarily for those vertices which are mapped to the nonzero real line. For any necklace $(v, \sigma(v), \sigma(\sigma(v)), \dots, \sigma^{-1}(v))$ mapped to the origin, simply break ties among $\epsilon(v), \epsilon(\sigma(v)), \epsilon(\sigma(\sigma(v))), \dots, \epsilon(\sigma^{-1}(v))$ so that the y -coordinates of the vertices in the necklace form a bitonic sequence. This completes the construction.

The techniques of Sections 5 and 6 may still be of value despite this improved layout. The results of Section 5 have applications beyond the area of VLSI layouts [4], and a better layout may yet be obtained using those methods. It is possible, however, that the layout presented in this addendum may be improved in other ways. For example, controlling the assignment of x -coordinates to necklaces might allow several exchange edges to occupy different parts of the same track.