

A Lower Bound on the Area of Permutation Layouts¹

Alok Aggarwal,² Maria Klawe,³ David Lichtenstein,⁴ Nathan Linial,⁵
and Avi Wigderson⁵

Abstract. In this paper we prove a tight $\Omega(n^3)$ lower bound on the area of rectilinear grids which allow a permutation layout of n inputs and n outputs. Previously, the best lower bound for the area of permutation layouts with arbitrary placement of the inputs and outputs was $\Omega(n^2)$, though Cutler and Shiloach [CS] proved an $\Omega(n^{2.5})$ lower bound for permutation layouts in which the set of inputs and the set of outputs each lie on horizontal lines. Our lower bound also holds for permutation layouts in multilayer grids as long as a fixed fraction of the paths do not change layers.

Key Words. VLSI, Grid embedding, Lower bound, Permutation, Layout, Multilayer.

1. Introduction. A rectangle R of a rectilinear grid is said to contain an n -permutation layout if $2n$ distinct nodes of R can be labeled as n inputs and n outputs so that for any one-to-one mapping τ from a subset of the inputs into the set of outputs, we can find a set of node disjoint paths inside R joining each input x in the domain of τ to the output $\tau(x)$. We refer to τ as a permutation, and to the node disjoint paths as the routing of τ . Figure 1 shows an example of such a routing.

Permutation layouts have proved extremely useful in the layout of printed circuits and large-scale integrated chips, and hence have received substantial attention (see [CS], [KKF], [S2], and [TK]). Examples of rectangles with $O(n^3)$ area which contain n -permutation layouts (and simple algorithms for finding the routings) were given by Cutler and Shiloach [CS], who also proved that if all the inputs lie on a single horizontal line, and all the outputs on another, then the rectangle must have area at least $\Omega(n^{2.5})$.

In this paper we prove that Cutler and Shiloach's upper bound is asymptotically optimal by giving an $\Omega(n^3)$ lower bound on the area of rectangles containing n -permutation layouts (a sketch of the proof of this result appeared in our conference paper [AKLLW]). Here we actually prove a stronger result, namely that multilayer grids which contain n -permutation layouts also require $\Omega(n^3)$ area,

¹ The first author's research was partially supported by NSF Grant No. MCS 820-5167. The third author's research was supported by NSF Grant No. MCS-8204246 and by a Hebrew University Fellowship.

² IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598, USA.

³ IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120-6099, USA. Current address: Department of Computer Science, University of British Columbia, Vancouver, BC, Canada V6T 1W5.

⁴ Yale University, New Haven, Connecticut, USA.

⁵ Hebrew University of Jerusalem, Jerusalem, Israel.




Fig. 1. A routing of a permutation.

as long as some (arbitrarily small) fixed fraction of the paths in each routing do not change layers. The restriction that a fixed fraction of the routing paths do not change layers is essential, since the standard crosspoint switch (see Figure 2) is a two layer n -permutation layout with $O(n^2)$ area, in which every routing path changes layers once. This area is asymptotically optimal (Theorem 5.2).

In the crosspoint switch it is easy to see that the grid nodes at which the path layer changes occur depend on which permutation is being routed. Paterson [P] gives a two-layer grid with $O(n^{2.5})$ area containing an n -permutation layout in which all the paths in the routing change layers exactly once, and these layer changes occur at a fixed set of n grid nodes, independent of the permutation being routed. Paterson's result is the best upper bound known currently for an n -permutation layout in which each path changes layers at most once, with all layer changes occurring at some fixed set of $O(n)$ nodes. However, as was pointed by our referee, it is possible to achieve an $O(n^2)$ area 2-layer n -permutation network, in which every path changes layers exactly twice with layer changes occurring at a fixed set of $2n$ nodes, by using (essentially) a $(\sqrt{n}, \sqrt{n}, \sqrt{n})$ Beneš network [B]. We sketch this construction at the end of Section 5.

The results in this paper imply lower bounds on the area needed for embedding n -node graphs of bounded degree in the k active layer model, a model which represents the layout of printed circuit boards with k layers. Upper bounds for this and related problems are given in our conference paper [AKLLW] and in [AKS].

The paper is organized as follows. In Section 2 we prove the existence of graphs with the property that any induced subgraph with enough edges has a quadratic crossing number. Section 3 defines a kind of quotient graph, and proves a straightforward lemma relating the crossing numbers of graphs and their quotients. In Section 4 we begin with a generalization of the definition of an n -permutation




Fig. 2. A crosspoint switch.

layout to multilayer grids. We then prove a few simple lemmas about rectangles containing multilayer permutation layouts. Finally, in Section 5 we prove the lower bound on the area of a rectangle containing a multilayer permutation layout, by examining a routing which yields a quotient graph with the crossing number properties discussed in Section 2. The lower bounds on the crossing number of the quotient graph and the relationship of the quotient graph's crossing number to the crossing number of the routing graph produce the desired lower bound on the area of the rectangle.

2. Crossing Numbers and Expanding Graphs. The key point in the proof of the lower bound is to use a routing to simulate a graph with a large crossing number. For the argument to work, it is necessary that the graph with a large crossing number have bounded degree. It is not hard to show that n -node expanding graphs have an $\Omega(n^2)$ crossing number (for bounded degree expanding graphs this follows immediately from Theorem 1 of [L2], and a simple extension to expanding graphs of arbitrary degree is given in Lemma 4.4 of [AKLLW]). Since the existence of expanding graphs of bounded degree is well known, in [AKLLW] we were able to use this lower bound to prove a lower bound on the area of rectangles containing n -permutation layouts. In multilayer permutation layouts, many of the paths in the routing will be allowed to change layers (a more precise definition is given in the next section), with only a fraction of the paths being forced to lie on a single layer. Thus, in order to prove a lower bound on the area of rectangles containing multilayer permutation layouts, we will need graphs of bounded degree with the property that every subset of sufficiently many edges induces a graph with a large

crossing number. It turns out that the right graphs for this purpose are regular graphs whose adjacency matrices have small enough second-largest eigenvalues. For regular graphs this eigenvalue property implies expansion; in fact, it is almost equivalent to being an expanding graph [AM], [A].

A graph is said to be d -regular if every node has degree exactly d . A graph is bipartite if the nodes of the graph can be partitioned into two disjoint sets, A and B , so that every edge is of the form (a, b) , where $a \in A$ and $b \in B$. It is easy to see that if a bipartite graph is also d -regular, then the sets A and B must have the same number of nodes.

LEMMA 2.1. *For each positive integer s , there is a positive integer $d(s)$, such that if n is sufficiently large, for some m with $n/2 < m \leq n$, there is a $d(s)$ -regular bipartite graph G on $2m$ nodes, with the property that, for any subset T of at most $\lceil 2m/s \rceil$ nodes, the number of edges of G with both endpoints in T is at most $8d(s)m/s^2$.*

The proof relies on two results concerning the second-largest eigenvalue of the adjacency matrix of a regular graph.

PROPOSITION 2.2 [AC, Proposition 4], [FP, Lemma 2.1]. *If λ is the second-largest eigenvalue of the adjacency matrix of a d -regular graph on m nodes, and S is a subset of αm nodes, then the number of edges with both endpoints in S is at most $\alpha^2 dm/2 + \lambda\alpha(1 - \alpha)m/2$.*

THEOREM 2.3 [LPS]. *If r and q are primes congruent to 1 mod 4 with $r < q$, there is an $(r + 1)$ -regular graph $X^{r, q}$ on $q + 1$ nodes, such that the second-largest eigenvalue, λ , of the adjacency matrix of $X^{r, q}$ satisfies $\lambda \leq 2\sqrt{r}$.*

PROOF OF LEMMA 2.1. Let r be the smallest prime congruent to 1 mod 4 with $r \geq s^2$. By the prime number theorem for arithmetic progressions [D, p. 123], for n sufficiently large there is a prime q with $r < n/2 < q < n$ which is congruent to 1 mod 4. Let G be the double cover of the $(r + 1)$ -regular $(q + 1)$ -node graph $X^{r, q}$, i.e., G has nodes $\{u_x : x \in X^{r, q}\} \cup \{v_x : x \in X^{r, q}\}$ and edges $\{(u_x, v_y) : (x, y) \text{ is an edge in } X^{r, q}\}$. We take $d(s) = r + 1$ and $m = q + 1$. Now let T be a subset of at most $\lceil 2m/s \rceil$ of the nodes of G , and let S be the subset of nodes of $X^{r, q}$ defined by $S = \{x : u_x \in T \text{ or } v_x \in T\}$. Clearly, $|S| \leq |T| \leq \lceil 2|X^{r, q}|/s \rceil$. Applying Proposition 2.2, a routine calculation shows that the number of edges of $X^{r, q}$ with both endpoints in S is at most $4d(s)m/s^2$. Finally, it is easy to see that the number of edges in G with both endpoints in T is at most twice the number of edges in $X^{r, q}$ with both endpoints in S , which completes the proof. \square

We are now ready to prove the result we need. The proof idea is essentially a variant on that of Theorem 1 in [L2]. We are grateful to Tom Leighton for suggesting this simplification of our original proof.

THEOREM 2.4. *For each positive integer p , there is a positive integer $d'(p)$, such that if n is sufficiently large, for some m with $n/2 < m \leq n$, there is a $d'(p)$ -regular bipartite*

graph G on $2m$ nodes, with the property that any subset of at least $1/p$ of the edges of G induces a subgraph with an $\Omega(n^2)$ crossing number.

PROOF. Let $s = 16p$ and assume that n is sufficiently large for Lemma 2.1 to hold. Define $d'(p) = d(s)$. Let G be the $d'(p)$ -regular bipartite graph on $2m$ nodes described in Lemma 2.1. Note that G has $d(s)m$ edges. Assume F is a subset of the edges of G with $|F| \geq d(s)m/p$, and let H be the subgraph of G induced by the edges in F . Next let c be the crossing number of H , and let b be the number of edges which must be removed from H to separate H into s pieces each of size at most $\lceil 2m/s \rceil$. Since $m > n/2$, to show that $c = \Omega(n^2)$, it suffices to prove that $c + 2m = \Omega(b^2)$ and $b = \Omega(m)$.

We first show that $b \geq d(s)m/2p$. Suppose B is a subset of F such that removal of the edges in B separates H into s disconnected pieces of size at most $\lceil 2m/s \rceil$. It suffices to show that $|B| \geq d(s)m/2p$. Let T_i be the set of nodes in the i th piece for $i = 1, \dots, s$. By the choice of G , the number of edges of F with both points in T_i is at most $8d(s)m/s^2$ for each i , so the total number of edges of F which are not in B is at most $8d(s)m/s = d(s)m/2p$. Now as F has at least $d(s)m/p$ edges, obviously $|B| \geq d(s)m/2p$.

Next, given a planar embedding of H with c edge crossings, we construct a $(c + 2m)$ -node planar graph, H' , by replacing each edge crossing with a new node. Assigning weight 1 to each original node of H , and weight 0 to each new node, by the weighted version of the planar separator theorem [LT1], there is a subset C of nodes of H' with $c + 2m = \Omega(|C|^2)$, whose removal separates the remaining nodes of H' into s pieces, each of weight at most $\lceil 2m/s \rceil$. Let B be the edges in F which are either adjacent to original nodes in C , or have edge crossings which are new nodes in C . Clearly $|C| = \Omega(|B|)$ and $|B| \geq b$; hence $c + 2m = \Omega(b^2)$. \square

It is more convenient for us to use this theorem in the following form.

COROLLARY 2.5. *For p fixed, there exist fixed $f, \xi > 0$, such that for every sufficiently large n there is an f -regular bipartite graph G on at most $2n$ nodes, such that any subset of at least $1/p$ of the edges of G induces a subgraph with at least ξn^2 crossing number.*

3. Quotient Graphs and Crossing Number. The purpose of this section is to introduce concepts which we use in Section 5 to formalize our notion of how a routed permutation “simulates” a graph of bounded degree, and in particular to demonstrate the connection between the crossing number of a graph and its “simulation.” For any graph G , a path partition of G is a collection $\{P_i : 1 \leq i \leq m\}$ of node disjoint simple paths in G such that every node of G is on some path P_i . If $Q = \{P_i : 1 \leq i \leq m\}$ is a path partition of G , the quotient graph of G by Q is the graph on m nodes $\{p_i : 1 \leq i \leq m\}$, where p_i is adjacent to p_j if and only if $i \neq j$ and there is some node on P_i adjacent to a node on P_j . When we apply this in Section 5, the quotient graph will be the graph we wish to simulate, and the original graph G will be a graph composed of permutation edges and some extra paths which form the path partition.

LEMMA 3.1. *Let $Q = \{P_i : 1 \leq i \leq m\}$ be a path partition of a graph $G = (V, E)$, let H be the quotient graph of G by Q , let d_H be the maximum degree of H , and let $E_Q = \{(x, y) \in E : (x, y) \text{ is an edge of } P_i \text{ for some } i\}$. Then for any embedding of G in the plane with the property that every edge crossing is between an edge in E_Q and an edge in $E \setminus E_Q$, the number of crossings in this embedding is at least $1/d_H$ times the crossing number of H .*

PROOF. It suffices to show that given such an embedding of G we can obtain an embedding of H in the plane with at most d_H times as many crossings. For each path P_i choose one of its endpoints to be p_i , and embed p_i wherever that endpoint of P_i was embedded. For each edge (p_i, p_j) in H , we choose an edge (x, y) joining a node x on P_i to y on P_j . The edge (p_i, p_j) is embedded by approximately following the embedding of P_i from p_i to x , then the embedding of (x, y) to y , and finally the embedding of P_j to p_j . It is not hard to see that this can be done so that whenever (x_j, y_j) is an edge of P_j and (x, y) is an edge of G which crosses (x_j, y_j) in the embedding of G , then at most d crossings occur between the edges of H adjacent to p_j and an edge of H following the embedding of (x, y) , where d is the degree of p_j . Moreover, the embedding can be done so that these are the only kind of edge crossings in the embedding of H . An example is shown in Figure 3. □




Fig. 3. Embedding the quotient graph.

4. Multilayer Permutation Layouts. A grid placement of a set V of nodes is a pair (σ, R) , where R is a rectangular grid and σ is a one-to-one mapping of V into the nodes of R . For k a positive integer, let R^{*k} be the three-dimensional rectilinear grid consisting of k copies of R aligned one on top of each other. Similarly, for z a node in R , let z^{*k} be the k vertically aligned nodes of R^{*k} which occupy z 's position in each copy of R .

If $G = (V, E)$ is a graph and (σ, R) is a grid placement of V , we say G has a k active layer (σ, R) embedding if there exists a set of paths $\{P_{(x, y)} : (x, y) \in E\}$ in R^{*k} which are node-disjoint except possibly at endpoints, with $P_{(x, y)}$ joining a node in $\sigma(x)^{*k}$ to a node in $\sigma(y)^{*k}$ for each edge (x, y) of G . We refer to the paths $P_{(x, y)}$ as the embedded edges of G , and sometimes, when the context is clear, as the edges of G . We say that a path in R^{*k} does not change layers if it remains entirely within one copy of R . Otherwise we say that the path changes layers. Examples of paths which do and do not change layers are shown in Figure 4.

For $0 \leq \alpha \leq 1$, we say G has a k active layer (σ, R) embedding with cut ratio α , if G can be embedded in R^{*k} with the placement of the nodes defined by σ , with at most $\alpha|E|$ of the embedded edges changing layers.

A rectangle R is said to contain a k -layer α -cut n -permutation layout if there is a grid placement (σ, R) of n inputs and n outputs such that, for every one-to-one mapping τ from a subset of the inputs into the outputs, the graph with edges $\{(x, \tau(x)) : x \text{ is an input in the domain of } \tau\}$ has a k active layer (σ, R) embedding with cut ratio α . We refer to such a grid placement as k -layer α -cut permutable, to the one-to-one mappings from subsets of the inputs to subsets of the outputs as permutations, and to the embedded edges of a permutation as the routing of the permutation. Moreover, we abbreviate k -layer 1-cut n -permutation layout to k -layer n -permutation layout. Note that a 1-layer n -permutation layout is simply the n -permutation layout we defined in the introduction.

In the next section we prove that, for each $k \geq 1$ and $\alpha < 1$, if n is sufficiently large, every rectangle which contains a k -layer α -cut n -permutation layout, has $\Omega(n^3)$ area. The remainder of this section is devoted to some simple lemmas about rectangles containing k -layer n -permutation layouts.

A horizontal (vertical) midline is any line halfway between two adjacent horizontal (vertical) lines of the grid.




Fig. 4. Path (b) changes layers, but path (a) does not.

LEMMA 4.1 *Let R be a rectangle containing x inputs and y outputs. Then either R has a horizontal line L containing at least $x/3$ inputs and $y/3$ outputs, or R has a horizontal midline L such that at least $x/3$ inputs lie on one side of L and $y/3$ outputs on the other.*

PROOF. Let L be the lowest horizontal midline such that L has either at least $x/3$ inputs or at least $y/3$ outputs below it. Without loss of generality assume that L has at least $x/3$ inputs below it. If L has at least $y/3$ outputs above it we are done so we may assume that there are less than $y/3$ outputs above L , and hence more than $2y/3$ outputs below L . Now if there are at least $x/3$ inputs above L we are done so we may assume there are greater than $2x/3$ inputs below L . Let L' be the horizontal line immediately below L . By the choice of L , L' has less than $x/3$ inputs and less than $y/3$ outputs strictly below it, and hence must have at least $x/3$ inputs and $y/3$ outputs on it. \square

COROLLARY 4.2. *If R contains a k -layer n -permutation layout, then R has both width and height at least $n/3k$.*

PROOF. If some horizontal line has at least $n/3$ inputs and $n/3$ outputs on it, then the width of R is at least $2n/3$. If not, then by Lemma 4.1 there is a horizontal midline L with at least $n/3$ inputs on one side and at least $n/3$ outputs on the other. Now any embedding of a permutation mapping $n/3$ of these inputs to $n/3$ of these outputs must have at least $n/3$ paths crossing L , and on each layer the paths must be disjoint. Thus in some layer there are at least $n/3k$ disjoint paths crossing L , so L (hence R) has width at least $n/3k$. A symmetric argument yields the same lower bound for the height. \square

COROLLARY 4.3. *If a rectangle R contains x inputs and y outputs, then there is a midline L (either horizontal or vertical) with $x/6$ inputs on one side and $y/6$ outputs on the other.*

PROOF. If there is a horizontal line L with at least $x/3$ inputs and at least $y/3$ outputs on it, let L' be the leftmost vertical midline such that, on its left, L' has either at least $x/6$ of the inputs on L or at least $y/6$ of the outputs on L . Clearly, L' must have either exactly $x/6$ inputs of L on its left or exactly $y/6$ outputs of L on its left but not both. Without loss of generality assume that exactly $x/6$ inputs are on the left of L' . Then since at most $y/6 - 1$ outputs of L are on the left of L' , at least $y/6$ outputs of L are on the right of L' . On the other hand, if no such horizontal line L exists, then by Lemma 4.1 there is a horizontal midline with at least $n/3$ inputs on one side and $n/3$ outputs on the other, so we are done. \square

LEMMA 4.4. *Let R be a rectangle of width W containing n inputs and n outputs, and let b be a positive integer. Then there are two collections I and O of subrectangles of R such that:*

- (a) *Each subrectangle has width at most b .*
- (b) *Any pair of distinct subrectangles in $I \cup O$ are disjoint.*




Fig. 6. Grouping the nodes into supernodes.

entire descending path must have very large height since its height is the sum of the heights of its subpaths. Thus our strategy is to find such a descending path with many long subpaths.

We do this by finding a “good” set of disjoint narrow subrectangles, and in each subrectangle taking either the descending path of minimal length which connects up all the inputs or the one connecting up all the outputs. This is shown in Figure 6(a). Along each path we group the inputs (outputs) into “supernodes” (see Figure 6(b)), and take the subpaths to be the subpaths connecting up the inputs (outputs) in each supernode. Then we choose a permutation so that the supernodes simulate a graph with a quadratic crossing number. Because of the quadratic crossing number, when this permutation is routed, many of the supernodes must have a linear number of the routed permutation edges crossing the subpath inside the supernode. A supernode subpath cannot be shorter than the number of routed permutation edges which cross it, so many of these supernode subpaths must be long. Finally, there must be paths which have lots of long supernode subpaths, and hence are very long. Of course, some amount of care is needed in the definitions of “narrow,” “many,” “long,” etc., to make this work, and multiple layers introduce some complications, but this is basically the idea behind the proof.

THEOREM 5.1. *For each $k \geq 1$ and $\alpha < 1$, if n is sufficiently large, every rectangle R containing a k -layer α -cut n -permutation layout has $\Omega(n^3)$ area.*

PROOF. Let (σ, R) be a grid placement of n inputs and n outputs which is k -layer permutable with cut ratio α . Choose an integer p such that $(1 - \alpha)/k \geq 1/p$, let f and ξ be as in Corollary 2.5, let $b = \lfloor \xi n/96f^3 \rfloor$, and let W be the width of R . We assume that n is large enough that $b \geq 1$. If $\lceil W/b \rceil \geq n/96f$ we are done since the height is at least $n/3k$ by Lemma 4.1 and $b = \Omega(n)$. Thus we may assume that $\lceil W/b \rceil < n/96f$.

Let I and O be collections of subrectangles of R with properties (a)–(e) of Lemma 4.4. Now for each rectangle in I , connect the inputs with a descending path P of shortest length on the grid, and group the inputs into sets of size f (f -tuples) according to their order on the path P . We refer to these f -tuples as superinputs. For each superinput let the height of the superinput be the vertical difference between the highest and lowest input in the superinput, and let the spine of the superinput be the subpath of P connecting the f inputs in the superinput. In each rectangle in I remove the $\lceil n/24fd \rceil$ superinputs with the greatest heights, where $d = \lceil W/b \rceil$.

We now prove that the total number of remaining superinputs is at least $\lceil n/48f \rceil$. Since $|I| \leq d$ and the rectangles in I collectively contained at least $n/12$ inputs, the number of superinputs originally created was at least $(n/12f) - d$. At most $d\lceil n/24fd \rceil$ superinputs were removed altogether so at least $\lceil (n/12f) - d - d\lceil n/24fd \rceil \rceil$ remain. Moreover, we assumed $d = \lceil W/b \rceil < n/96f$, so we have $\lceil (n/12f) - d - d\lceil n/24fd \rceil \rceil \geq \lceil (n/12f) - d - d(1 + n/24fd) \rceil = \lceil (n/24f) - 2d \rceil \geq \lceil n/48f \rceil$ as desired.

Similarly from superoutputs, remove the $\lceil n/24fd \rceil$ superoutputs with the greatest heights from each rectangle in O , leaving at least $\lceil n/48f \rceil$ superoutputs.

By Corollary 2.5, for n sufficiently large, there is an f -regular bipartite graph G on at most $2\lceil n/48f \rceil$ nodes such that any subset of at least $1/p$ of the edges of G induces a subgraph with at least an $\xi\lceil n/48f \rceil^2$ crossing number. Let $2m$ be the number of nodes in G , let E be the edges of G , and let A and B be the two node subsets of G , i.e., both $|A|, |B| = m$ and all edges in E are between A and B . Let V_I be a subset of exactly m of the remaining superinputs in I , and similarly let V_O be m of the remaining superoutputs in O . We refer to the superinputs and superoutputs in $V_I \cup V_O$ as supernodes.

We are ready to choose the permutation so that the supernodes simulate G . Let φ be a one-to-one mapping from $A \cup B$ onto $V_I \cup V_O$ with $\varphi(A) = V_I$ and $\varphi(B) = V_O$. Let $V'_x = \{z : z \text{ is an input in a superinput in } V_I\}$ and let $V'_o = \{z : z \text{ is an output of a superoutput in } V_O\}$. Now let τ be a permutation from V'_I to V'_O such that, for every edge (x, y) in E , there is an input x' in $\varphi(x)$ and an output y' in $\varphi(y)$ such that $y' = \tau(x')$. Since G is f -regular it is easy to see that such a τ exists. For each edge $(x, y) \in E$, let $\Phi(x, y)$ be the edge (x', y') , where x', y' are as above.

Since (σ, R) is a k -layer permutable grid placement with cut ratio α , there is a k -layer embedding, η , of the set of edges $\Phi(E)$ such that at most $\alpha|E|$ edges change layers. Since at least $(1 - \alpha)|E|$ edges of $\Phi(E)$ do not change layers, some layer must contain at least $(1 - \alpha)|E|/k$ edges of $\Phi(E)$ which make no layer changes. Assume

without loss of generality that this is the first layer, let $F \subseteq E$ such that $\Phi(F)$ is this set of edges, and let H be the subgraph of G induced by the edges in F . Note that since $|F| \geq (1 - \alpha)|E|/k \geq |E|/p$, by the choice of G the subgraph H has crossing number at least $\xi \lceil n/48f \rceil^2$.

Let $V' = V'_I \cup V'_O$, let $E' = \Phi(F) \cup \{(v', w') : v' \text{ and } w' \text{ are in the same supernode and are adjacent on its spine}\}$, and let $H' = (V', E')$. It is easy to see that the spines of the supernodes form a path partition of H' , and that H is the quotient graph of H' by this path partition. Moreover, the embedding of H' obtained by combining the embedding of $\Phi(F)$ given by η with the spine paths themselves, has the property that every edge crossing is between an edge in $\Phi(F)$ and a spine edge. Thus by Lemma 3.1 this embedding of H' has at least $\xi \lceil n/48f \rceil^2 / f$ crossings. Since there are $2m$ supernode spines and $m \leq \lceil n/48f \rceil$, one of these spines must have at least $\xi \lceil n/48f \rceil / 2f$ crossings, and hence must be at least that long. Since the original path was descending of shortest length, the spine must also be descending of shortest length, and it is easy to see that this implies that the supernode has height at least $(\xi \lceil n/48f \rceil / 2f) - (f - 1)b$. This is because the rectangle containing the supernode has width at most b , and clearly any f nodes can be joined by a descending path with length at most $(f - 1)b$ plus the maximum difference in their heights. Moreover, $(\xi \lceil n/48f \rceil / 2f) - (f - 1)b \geq \xi n / 96f^3$ since $b = \lfloor \xi n / 96f^3 \rfloor$.

Now, at least $\lceil n/24fd \rceil$ supernodes of greater height were removed from the rectangle containing this supernode. Thus, since the original path connecting the supernodes in this rectangle was descending, its total height is at least $\lceil n/24fd \rceil (\xi n / 96f^3)$. This implies that the area of R is at least $W \lceil n/24fd \rceil (\xi n / 96f^3)$, which is easily verified to be $\Omega(n^3)$ \square

THEOREM 5.2. *A similar (though much simpler) argument can be used to prove that, for any fixed $k \geq 1$, if n is sufficiently large, every rectangle containing a k -layer n -permutation layout has area $\Omega(n^2/k^2)$. Briefly, the argument is as follows. Since the n -permutation layout can simulate an expanding graph with a linear number of nodes, and since the expanding graph has a quadratic crossing number, when the simulated expanding graph is projected onto the bottom layer there must be $\Omega(n^2)$ crossings. It is not hard to see that this implies that one of the grid layers must have $\Omega(n^2/k^2)$ area.*

We close with a brief sketch of how the Beneš network can be used to give an $O(n^2)$ area 2-layer n -permutation network, in which every path changes layers exactly twice with layer changes occurring at a fixed set of $2n$ nodes. This idea was suggested by the referee. The Beneš network is a three-stage network, with each stage consisting of \sqrt{n} disjoint $\sqrt{n} \times \sqrt{n}$ crossbars. For $i = 1, 2$, the j th output in the k th crossbar of stage i is connected to the k th input in the j th crossbar of stage $(i + 1)$. It is well known that this forms an n -permutation layout between the n inputs of the first stage and the n outputs of the third stage. We now describe the 2-layer grid implementation of the Beneš network. An illustration of the layout is given in Figure 7. Each $\sqrt{n} \times \sqrt{n}$ crossbar is implemented by a \sqrt{n} -permutation layout in an $O(\sqrt{n}) \times O(n)$ 1-layer grid, in which the inputs lie in order on the top horizontal line, and the outputs are spaced $O(\sqrt{n})$ apart in order on the middle




Fig. 7. The 2-layer layout of the Beneš construction.

vertical line. (Such a layout is easy to construct using the methods of Cutler and Shiloach, and is the type illustrated in Figure 1.) Thus each stage can be embedded in a single layer in an $O(n) \times O(n)$ grid with all inputs in order on the top line. Moreover, because of the spacing of the outputs, it is easy to see that all the edges from the outputs of stage i to the inputs of stage $i + 1$ can be routed in a single layer in an $O(n) \times O(n)$ grid with the inputs of stage $i + 1$ lying in order on the bottom horizontal line. From this it is clear that the Beneš network can be embedded as follows. Stage 1 is embedded on layer 1. On layer 2 we embed the edges from the outputs of stage 1 to the inputs of stage 2 and stage 2 itself. Finally, back on layer 1, we embed the edges from the outputs of stage 2 to the inputs of stage 3 and stage 3. When used as an n permutation layout, each path changes from layer 1 to 2 at a stage 1 output and back to layer 1 at a stage 2 output.

Acknowledgements. We would like to thank Noga Alon, Ashok Chandra, Ruth Cohen, Danny Dolev, Rao Kosajaru, Tom Leighton, Gary Miller, Nick Pippenger, and Michael Rabin for comments and suggestions. In addition, we are especially grateful to Tom Leighton for his simplification of the proof of Theorem 2.4, and to the referee for the suggestion of the Beneš network construction.

Note Added in Proof. Klawe and Leighton have recently extended the $\Omega(n^3)$ lower bound on the size of permutation layouts to arbitrary planar graphs implementing a permutation network (see M. Klawe and Tom Leighton, A tight lower bound on the size of planar permutation networks, *Proc. Internat. Symp. SIGAL '90*, Lecture Notes in Computer Science, Vol. 450, Springer-Verlag, Berlin, pp. 281–287).

References

- [A] N. Alon, Eigenvalues and expanders, Preprint.
- [AM] N. Alon and V. Milman, λ_1 , isoperimetric inequalities for graphs, and superconcentrators, *J. Combin. Ser. Theory B*, **38** (1985), 73–88.
- [AC] N. Alon and F. R. K. Chung, Explicit Constructions of Linear-Sized Tolerant Networks, Preprint.
- [AKS] A. Aggarwal, M. Klawe, and P. Shor, Multilayer Grid Embeddings for VLSI, *Algorithmica*, **6** (1991), 129–151.
- [AKLWW] A. Aggarwal, M. Klawe, D. Lichtenstein, N. Linial, and A. Wigderson, Multi-layer grid embeddings, *Proc. IEEE 26th Ann. Symp. on Foundations of Computer Science*, 1985, pp. 186–196.
- [B] V. E. Beneš, Optimal rearrangeable multistage connecting networks, *Bell System Tech. J.*, **43** (1964), 1641–1656.
- [BP] N. K. Bose and K. A. Prabhu, Thickness of graphs with degree constrained vertices, *IEEE Trans. Circuits and Systems*, **24** (4) (1977), 184–190.
- [CS] M. Cutler and Y. Shiloach, Permutation layout, *Networks*, **8** (1978), 253–278.
- [D] H. Davenport, *Multiplicative Number Theory*, Markham, Chicago, 1967.
- [FP] J. Friedman and N. Pippenger, Expanding Graphs Contain All Small Trees, *Combinatorica*, **7** (1987), 71–76.
- [K] Y. Kajitani, On Via Hole Minimization of Routings on a Two-Layer Board, Tech. Report, Dept. of Elec. Engineering, Tokyo Institute of Technology, 1980.
- [KS] M. Klawe and P. Shor, Optimal Universal Layouts, Preprint, 1986.
- [KKF] E. S. Kuh, T. Kashiwabara, and T. Fujisawa, On optimum single row-routing, *IEEE Trans. Circuits and Systems*, **26** (6) (1979), 361–368.
- [L1] F. T. Leighton, Layouts for the Shuffle-Exchange Graph and Lower Bound Techniques for VLSI, Ph.D. Thesis, Dept. of Math., MIT, 1981.
- [L2] F. T. Leighton, New lower bound techniques for VLSI, *Math. Systems Theory*, **17** (1984), 47–70.
- [LPS] A. Lubotsky, R. Philips, and P. Sarnak, Explicit expanders and the Ramanujan conjectures, *Combinatorica*, **8** (1988), 261–278.
- [LT1] R. J. Lipton and R. E. Tarjan, A separator theorem for planar graphs, *SIAM J. Appl. Math.*, **36** (1979), 177–189.
- [LT2] R. J. Lipton and R. E. Tarjan, Applications of a planar separator theorem, *SIAM J. Comput.*, **9** (1980), 615–627.
- [MC] C. A. Mead and L. A. Conway, *Introduction to VLSI Systems*, Addison Wesley, Reading, Mass., 1980.
- [P] M. S. Paterson, Universal Chains and Wiring Layouts, Preprint.

- [R] R. Rivest, The placement and interconnect system, *Proc. 19th Design Automation Conf.* 1982, pp. 475-481.
- [S1] Y. Shiloach, Linear and Planar Arrangement of Graphs, Ph. D. Dissertation, Weizmann Institute, Rehovot, 1976.
- [S2] I. Shirakawa, Letter to the editor: some comments on permutation layout, *Networks*, **10** (1980) 179-182.
- [S3] H. C. So, Some theoretical results on the routing of Multilayer printed wiring boards, *Proc. IEEE Int. Symp. on Circuits and Systems*, 1974, pp. 296-303.
- [T] C. D. Thompson, A Complexity Theory for VLSI, Ph.D. Dissertation, Carnegie-Mellon University, 1980.
- [TKS] B. S. Ting, E. S. Kuh, and I. Shirakawa, The multilayer routing problem: algorithms and necessary and sufficient conditions for the single-row single layer case, *IEEE Trans. Circuits and Systems*, **23** (12) (1979), 768-778.
- [TK] S. Tsukiyama and E. S. Kuh, Double-row planar routing and permutation layout, *Networks*, **12** (1982), 287-316.
- [V] L. G. Valiant, Universality considerations in VLSI, *IEEE Trans. Comput.*, **30** (2) (1981) 135-140.