MULTI-LAYER GRID EMBEDDINGS

Alok Aggarwal, IBM T.J.Watson Research Center
Maria Klawe, IBM Almaden Research Center
David Lichtenstein, Yale University
Nathan Linial, Hebrew University of Jerusalem
Avi Wigderson, IBM San Jose Research Laboratory

ABSTRACT

In this paper we propose two new multi-layer grid models for VLSI layout, both of which take into account the number of contact cuts used. For the first model in which nodes "exist" only on one layer, we prove a tight area x (number of contact cuts) = $\Theta(n^2)$ trade-off for embedding any degree 4 n-node planar graph in two layers. For the second model in which nodes "exist" simultaneously on all layers, we prove a number of bounds on the area needed to embed graphs using no contact cuts. For example we prove that any n-node graph which is the union of two planar subgraphs can be embedded on two layers in $O(n^2)$ area without contact cuts. This bound is tight even if more layers and an unbounded number of contact cuts are allowed. We also show that planar graphs of bounded degree can be embedded on two layers in $O(n^{1.6})$ area without contact cuts.

These results use some interesting new results on embedding graphs in a single layer. In particular we give an $O(n^2)$ area embedding of planar graphs such that each edge makes a constant number of turns, and each exterior vertex has a path to the perimeter of the grid making a constant number of turns. We also prove a tight $\Omega(n^3)$ lower bound on the area of grid n-permutation networks.

1. INTRODUCTION

The problem of embedding bounded degree graphs in rectilinear grids has been studied extensively in recent years [Lei80, Th80, Va81]. In these papers, the grid graph in which a graph G is to be embedded consists of nodes on plane points with integer coordinates, and the edges of the grid join exactly those nodes which are unit distance apart. Consequently, an "embedded edge" of G consists of a path of horizontal and/or vertical segments such that a horizontal segment of an embedded edge can only intersect a vertical segment of some other embedded edge. Here the cost of the embedding is usually measured by either its area, maximum edge length, or the minimum number of crossings. This notion of embedding is particularly useful in the fabrication of VLSI chips and in the design of printed circuit boards. However, it ignores the important problem of layer changes. In practice, the wires that cross each other must be on different layers. Thus whenever they are on the same layer, one of them must change its layer before the crossing occurs. This layer change is achieved by making a "via" or a "contact cut" between the two layers. Usually, the presence of too many contact cuts leads to a larger area, requires the design of expensive masking strategies, and leads to a deterioration in performance with a higher probability of faulty chips. Hence, most design automation systems such as the Placement-Interconnect system [Ri82] perform the arduous task of minimizing the number of contact cuts in a given layout. These design automation systems separate the problem of layer assignment from the geometrical issue regarding area consumption and wire lengths. However, in this paper, we study the two issues together, and obtain some interesting bounds and tradeoffs.

This paper is divided into seven sections. In section 2, we introduce two models corresponding to VLSI chips and printed circuit boards, and discuss some basic properties and relations between these models. In section 3, we establish a tight area-cut tradeoff for embedding planar graphs in the first model.

Section 4 deals with two types of grid permutation networks, n-path permutation graphs and n-cycle permutation graphs. We demonstrate an $O(n^3)$ area grid n-cycle permutation graph similar to Cutler and Shiloach's $O(n^3)$ area grid n-path permutation graph [CS78]. Using a lemma on the crossing number of expanding graphs we prove a matching lower bound on the area of both types of grid permutation networks.

Section 5 contains new results on embedding planar graphs in the one layer grid (with no crossings allowed). The first is an $O(n^2)$ area embedding guaranteeing that each edge makes at most 6 turns, and each exterior vertex has a path to the perimeter of the grid making at most 2 turns. Applying this with results on permutation networks, we obtain planar graph embedding algorithms which respect a fixed placement of the nodes in the grid.

In section 6, the results of sections 4 and 5 are combined to prove results about the area needed to embed various classes of graphs in the second model. For two active layers we prove a tight $\Theta(n^2)$ bound on the area needed to embed n-node graphs of thickness 2. For $k \geq 3$ we show that n-node graphs of thickness k can be embedded in $O(n^5)$ area in k active layers, and also that n-node graphs of degree at most 2k can be embedded in $O(n^3)$ area in k active layers. Finally we show that n-node planar graphs of bounded degree have an $O(n^{1.6})$ area embedding in two active layers. The last section, 7, contains open problems.

II. MODELS

An embedding of a graph G in the (one layer) grid is a mapping of the nodes of G to nodes of the grid, and edges of G to paths in the grid. The paths representing edges are disjoint except for the necessary intersections at their endpoints. Since the grid is planar and every node in the grid has degree at most four, only planar graphs of degree at most four can be embedded in this manner. In order to embed a planar graph G of arbitrary degree we sometimes use a *rake embedding*, namely we embed each node of G as a horizontal line segment with all incident edges entering from below. An example is shown in Figure 1. It is easy to see that a rake embedding of a graph of degree at most 4 can easily be converted into an ordinary grid embedding without increasing the area by more than a constant factor.

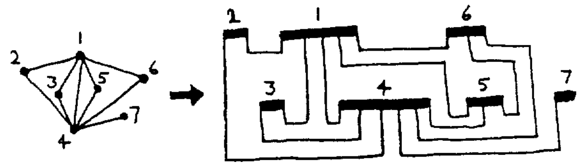

Figure 1. A rake embedding

The first multi-layer model, which we will call the *single active layer* or *SAL* model, consists of two grid layers which we will refer to as green and blue. All the nodes of G are embedded on the green layer, and edges are represented by paths in the grid which begin and end on the green layer, but may change layers any number of times at contact cuts. Within each layer no paths cross each other except for the obvious intersections at endpoints in the green layer. This model captures the present fabrication technology used in VLSI circuits, where the blue and green layers represent the metallic and polysilicon layers in a VLSI chip (for details see Mead and Conway [MC80]). Figure 2(a) shows an embedding of K_5 (the complete graph on five nodes) in this model. The blue and green segments of the edge paths are drawn as dashed and solid lines, respectively, and contact cuts are denoted by small squares.

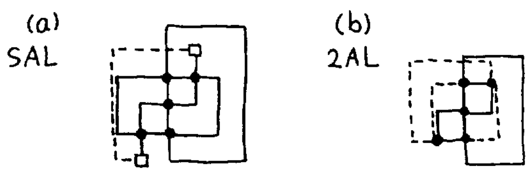

Figure 2. SAL and 2AL embeddings of K_5

Leiserson [Lei80] gave a simple technique for embedding any graph with m edges in the SAL model with $O(m^2)$ area and O(m) contact cuts. He also showed that n-node planar graphs have an $O(n \log^2 n)$ area embedding in this model. Leighton [Leig81] proved an $\Omega(n \log n)$ area lower bound for embedding planar graphs in the SAL model.

We now define the *k active layer* or *kAL* model, where k is a positive integer. This model consists of k grid layers, each node of G is embedded in the same position on each layer, and edges are embedded as paths in the grid which may change layers at contact cuts. An edge path may begin and end on any layer, but as before, within each layer the paths must not cross except at their endpoints. It is easy to see that a 1AL embedding is just a one layer grid embedding in the usual sense. The kAL model corresponds to k layer printed circuit boards in which the pins of a mounted chip are present on all layers. Figure 2(b) shows an embedding of K_5 in the 2AL model.

The *thickness* of a graph G is the minimum number k such that G is the union of k planar graphs. (Here, by "union of k planar graphs" we mean that the edges of G can be partitioned into k sets so that the graph induced by each set is planar.)

Observation 2.1. A graph can be embedded in the kAL model without contact cuts if and only if it has thickness k.

Proof. It is clear that the number of layers needed to embed a graph G is at least its thickness, and a simple homotopy argument proves the opposite direction. □

Some of our results on the kAL model can be applied to a variant model in which each node of the graph must be embedded in a specified position. This is referred to as a fixed placement model, and is of importance for printed circuit boards since often the designer has no say regarding the placement of the nodes.

We close this section with some comments on the differences between the SAL and kAL models. First, without contact cuts only planar graphs can be embedded in the SAL model whereas graphs of thickness k can be embedded in the kAL model. Moreover, it can be shown that some important interconnection networks such as shuffle exchange networks, cube connected cycles, and meshes of trees have an embedding without contact cuts in the 2AL model with no increase in area relative to the standard model. In contrast, shuffle exchange networks and cube connected cycles require $\Omega(n/\log^2 n)$ contact cuts in the SAL model, regardless of the area spent. There are also examples of degree 4, n-node graphs with constant size separators which require $\Omega(n)$ contacts cuts for any SAL embedding. This shows that separator size is not a useful concept for the SAL model.

3. PLANAR AREA-CUT TRADEOFF IN THE SAL MODEL

In this section we show that for each C with $1 \leq C \leq n/\log^2 n$, every n-node planar graph of degree at most 4 has an SAL embedding in $O(n^2/C)$ area with O(C) contact cuts. Moreover, for each n and $C \leq n/4$ there is an n-node planar graph for which every SAL embedding with C contact cuts has area $\Omega(n^2/C)$. This result can easily be generalized to planar graphs of bounded degree using rake embeddings. The upper bound is obtained using a hybrid algorithm, based on three different embedding techniques, namely [Lei80], [DLT81] and [LT80]. The lower bound extends the work of Shiloach [Sh76] and Valiant [Va81], which gives a bound for one extreme of the tradeoff spectrum, to the whole spectrum.

Before proving this section's main result we present the three embedding techniques and a useful lemma.

A planar graph is said to have *gauge* at most g if there is a planar embedding of the graph such that every node has a path of length at most g to a node on the outside face. Dolev, Leighton and Trickey prove the following result in [DLT84]. (In [DLT84] gauge is called width but in this paper we use the term width in a more conventional sense.)

Theorem 3.1. [DLT84] Every n-node bounded degree planar graph of gauge g has a (one layer) embedding in a square grid with $O((ng)^{1/2})$ side length, hence O(ng) area.

We will use the following lemma to subdivide the graph into subgraphs of small gauge. In fact we prove a slightly stronger version than needed here, as we will use the stronger version in section 6.

Lemma 3.2. Let G be an n-node planar graph of bounded degree and let $g \geq 1$. Then the edges of G can be partitioned into two sets, E_1 and E_2, such that the subgraphs, G_1 and G_2 induced by E_1 and E_2 have the following properties. The gauge of G_1 is at most g, the gauge of G_2 is at most 4, and $|E_2| = O(n/g)$.

Proof. Fix a planar embedding of G. Let V_i be the set of nodes in G whose shortest path to a node on the outside face has length i. For each $j = 1,\ldots,\lceil 2n/g \rceil$ find the V_i with $(j-1)g/2 \leq i < jg/2$ such that V_i has the smallest cardinality, and place all the edges incident with its nodes in E_2. $|E_1|$ is taken to be the set of edges which remain. It is easy to see that $|E_2| = O(n/g)$, that none of the connected components of $|G_1|$ have gauge more than g, and that none of the connected components of G_2 have gauge more than 4. The proof is completed by observing that the gauge of a graph is the maximum of the gauges of its connected components. □

The second embedding strategy, using a simple technique of Leiserson [Lei80], adds edges to a SAL embedding of a subgraph using a constant number of contact cuts per edge.

Lemma 3.3. [Lei80] Let G' be a subgraph of a degree 4 planar graph G and suppose G' has a SAL embedding in a square of side S such that all horizontal segments of edges lie on the green layer. Then e edges of G can be added to the embedding using at most 4e contact cuts so that the resulting embedding is contained in a square of side S + 3e. Moreover all horizontal segments of the added edges will also lie on the green layer.

The final strategy, also used by Leiserson in [Lei80], is a traditional divide and conquer technique based on iterating a variant of the planar separator theorem, then recombining the pieces using the preceding lemma.

Lemma 3.4. There is a constant k such that any degree 4, n-node planar graph can be partitioned into four disconnected subgraphs of size (approximately) n/4 by the removal of $kn^{1/2}$ edges.

Proof. This is easily deduced from the usual version of the planar separator theorem (Lipton and Tarjan [LT79]). □

Theorem 3.5. For any C with $1 \leq C \leq n/\log^2 n$, every n-node planar graph of degree at most 4 has an SAL embedding in $O(n^2/C)$ area with O(C) contact cuts. Moreover, for each n and $C \leq n/4$ there is an n-node planar graph for which every SAL embedding with C contact cuts has area $\Omega(n^2/C)$.

Proof. The proof of the lower bound goes as follows. Shiloach [Sh76] proved that the n-node planar graph shown in Figure 3 requires $\Omega(n^2)$ wire area if no contact cuts are allowed. We will call this graph a triangle graph. (Valiant [Va81] mentioned a similar result for a diamond shaped graph.) Given an embedding of the triangle graph with C contact cuts delete all edges in the graph which use a contact cut. The embedding of the remaining graph uses no contact cuts, and it is easy to see that the remaining graph contains C disjoint triangle subgraphs containing a total of at least n - 3C nodes. Now the total area for the remaining graph is at least the sum of the wire areas of the C triangle subgraphs, and it easy to see that this sum is minimized when all the triangle subgraphs have the same size, yielding a lower bound of approximately $C((n - 3C)/C)^2 = (n - 3C)^2/C = \Omega(n^2/C)$.

Due to space limitations, we only sketch the proof of the upper bound. Moreover we will omit floors and ceilings, treating all quantities that obviously need to be integers as though they already are. To prove the upper bound we show how to embed any degree 4 planar graph with O(C) contact cuts and $O(n^2/C)$ area. This construction mixes embedding strategies according to the number of contact cuts allowed.

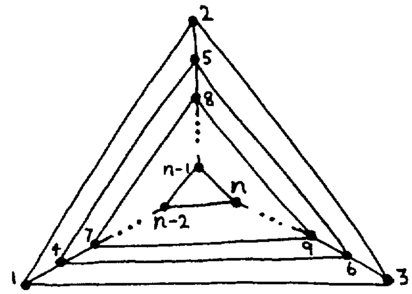

Figure 3. The n-node triangle graph

Case 1:
First suppose that $C \leq n^{2/3}$. Let $g = n/C$. Remove $O(n/g) = O(C)$ edges so that each connected component has gauge at most g. Use Theorem 3.1 to embed each component G' in the green layer in a square grid of side $O((n'g)^{1/2})$, where n' is the number of nodes in G'. Using a simple 2 dimensional bin packing algorithm, these "small" square grids can be packed into a "large" square grid whose area is at most a constant factor greater than the sum of the areas of the "small" square grids. Thus the large square has area O(ng) and hence side length $O((ng)^{1/2}) = O(n/C^{1/2})$. Now we have an embedding of all of G on the green layer, except for the O(C) edges we removed. Using the construction in Lemma 3.3 to embed these edges yields a rectangle of height and width $O(n/C^{1/2} + C)$, and hence area $O(n^2/C + C^2) = O(n^2/C)$ since $C \leq n^{2/3}$. Moreover all horizontal segments of edges lie on the green layer.

Case 2:
Let $n^{2/3} < C \leq n/(\log n)^2$. We recursively apply Lemma 3.4 j times where $4^j = C^3/n^2$. This process corresponds to constructing a complete quad tree of depth j, where the root is the original graph, and the four children of an internal node H are the four components into which H is partitioned by applying Lemma 3.4. Now we prove by backwards induction on i that there is a constant K, such that for any graph H of depth i in the tree, H has an SAL embedding in a square of side $2^{j-i}Kn^2/C^2 + (j-i)3kn^{1/2}/2^i$. The base case uses the embedding in case 1 with $O((n/4^j)^{2/3})$ contact cuts per leaf graph. For the inductive step we place the embeddings of the four children of H in a 2 x 2 square. Then using Lemma 3.3 we add in the $kn^{1/2}/2^i$ edges which were removed in partitioning H into its "children". Simple algebra verifies that at the root the total number of contact cuts used is O(C) and the area is $O(n^2/C)$ as long as $C \leq n/\log^2 n$. □

4. PERMUTATION NETWORKS

In this section we study one layer grid embeddings of two types of permutation networks. An *n-path permutation graph* is a graph with two sets of n distinguished nodes called inputs and outputs such that for every one-to-one mapping from the inputs to the outputs there is a set of node disjoint paths joining each input to the output specified by the mapping. We say that such a set of paths realizes the mapping. Similarly *an n-cycle permutation graph* is a graph with n distinguished nodes called terminals, such that for every permutation on the terminals, the graph contains a set of node disjoint cycles which realize the permutation. Cutler and Shiloach [CS78] showed that there is a grid with $O(n^3)$ area which is an n-path permutation graph. We include (essentially) their construction here as we will need to refer to its details in the next section. Using similar techniques we will also show that there is a grid with $O(n^3)$ area which is an n-cycle permutation graph. Finally using a lemma on the crossing number of expanding graphs we prove matching lower bounds, namely that any grid which is an n-permutation graph (of either variety) must have $\Omega(n^3)$ area. This improves on the result in [CS78] which proved a lower bound of $\Omega(n^{2.5})$ area for the special case of n-path permutation graphs which have all inputs on a single line and all outputs on a single line.

Proposition 4.1 (Cutler and Shiloach). The $n^2 \times 2n$ grid is an n-path permutation graph.

Proof. The inputs are placed on the top line in the first n positions. The outputs are placed on the middle vertical line at intervals of length n beginning at the (n+1)-st position. Given a one-to-one mapping σ from the inputs to the outputs the paths realizing σ are as follows. The paths leave the inputs forming a "ribbon" of width n. The ribbon immediately jogs right far enough so that the path from the input which σ maps to the first output is on the middle vertical line and thus can run down and connect to the first output. Between the first and second outputs, the remaining paths in the ribbon jog horizontally so that the path to be connected to the second output is on the middle vertical line and thus can be connected. The process continues with the ribbon making an appropriate horizontal jog after passing each output so that the next output can be connected to its path. Since the width of the ribbon is at most n at all times, there are enough horizontal lines between each pair of outputs to make the jog, and the ribbon always stays within the vertical boundaries of the rectangle. An example of such a set of paths is shown in Figure 4. □

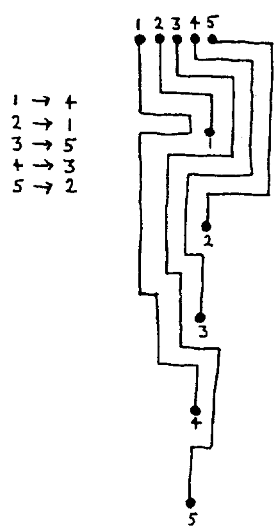

Figure 4. Routing paths in the 5-path permutation graph

Remark 4.2. Let us call a placement of n inputs and n outputs in a rectangle well-spaced if all the inputs lie on one horizontal line, all outputs lie at least n lines below the inputs and at least n lines away from both the vertical edges of the rectangle, and if there is a vertical gap of at least n between any pair of outputs. It is easy to see that the same idea as in 4.1 can be used to prove that any rectangular grid with a well-spaced placement of n inputs and n outputs is an n-path permutation graph.

Proposition 4.3. The $4n^2 \times 4n$ grid is an n-cycle permutation graph.

Proof. The terminals 1,2,...,n are placed in order down the middle vertical line at intervals of 4n. Clearly it suffices to show how to embed any set of disjoint cycles on 1,2,...,n given this placement. The general idea is that the cycles are embedded one at a time, and the cycle edges are embedded in the order they occur in the cycle, beginning with an edge leaving the topmost node in the cycle. At any given time the embedding is comb shaped, and those nodes whose cycle edges have not been embedded are exposed between teeth of the comb. The spine of the comb is initialized to be the rightmost vertical line of the grid. Routing a cycle edge (x,y) consists of running the edge horizontally right from x until reaching the current spine of the comb, then following the outline of the comb until level with y, then left to y (thus beginning a new tooth at y except when (x,y) is the last edge of the cycle). In the case that x = y, i.e. the edge is a loop, after running across to the spine the edge runs vertically for one unit then back across to the middle line and into x. Each cycle edge uses at most 2 horizontal lines between any pair of adjacent nodes so the 4n spacing between nodes suffices. Moreover for any cycle edge which is not a loop, none of its vertical segments overlap. From this it is easy to see that at most 2n vertical lines on either side of the middle vertical line are needed. □

Figure 5(a) shows the embedding of the cycle (1,5,3,8) and Figure 5(b) shows how the cycles (2,7,4) and (6) are added to this embedding.

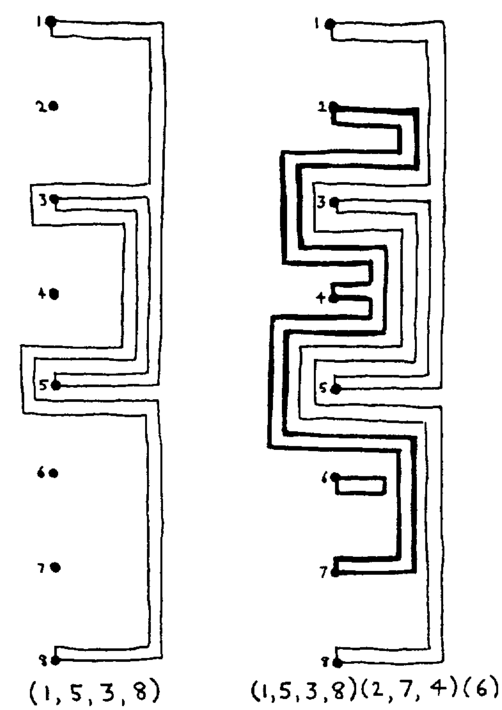

Figure 5. Routing cycles in the 8-cycle permutation graph

We now prove a lemma on the crossing number of an expanding graph. This lemma will be a key ingredient in the lower bounds for permutation graphs and for embedding graphs in the kAL models. Let $d > 0$. We say that an n-node graph is a d-expanding graph if every subset A of at most n/2 nodes is adjacent to at least $d|A|$ nodes outside A. It is well known that for any $d < 1$ there exist d-expanding n-node graphs of bounded degree for any n (see [C78] for example). For example, if n is even, a random bipartite graph of degree 3 between two sets of n/2 nodes is almost always an expanding graph. In particular, [M] proves the existence of a 1/6-expanding graph of degree 3.

Lemma 4.4. Every embedding of an n-node expanding graph in the plane with the property that at most two edges intersect at any point which is not a node, must have $\Omega(n^2)$ edge crossings.

Proof. Let G be an n-node d-expanding graph. Given an embedding of G in the plane with the above property, we construct a planar graph G' by placing a new node at each edge crossing. Let n' be the total number of nodes in G'. We assign each original node of G the weight 1, and the weight 0 to all the new nodes. By the weighted planar separator theorem [LT79, Corollary 3] there is a set C of nodes which separates G' into two sets A and B, such that both A and B have total weight at most n/2, and $n' = \Omega(|C|^2)$. It suffices to show that C has at least dn/(10+d) nodes since $|C| \geq dn/(10+d)$ implies $n' = \Omega(n^2)$, and hence the number of edge crossings, n' - n, is also $\Omega(n^2)$. Suppose $|C| < dn/(10+d)$. Then the weight of C is less than dn/(10+d) so one of A and B, say A, has weight at least n(1 - d/(10+d))/2 and hence at least this many original nodes of G. Since G is d-expanding, there are at least dn(1 - d/(10+d))/2 original nodes in $B \cup C$ which are adjacent in G to nodes in A. By the assumption about the size of C, at least dn(1 - d/(10+d) - 2/(10+d))/2 = 4dn/(10+d) of these are in B. Thus there are 4dn/(10+d) edge disjoint paths from A to B in G'. As C separates A from B there are at least this many edges adjacent to new nodes which are in C. Finally each new node has degree 4 so that are at least dn/(10+d) new nodes in C, a contradiction. □

Theorem 4.5. Every n-cycle permutation graph which is a subgraph of the grid has area $\Omega(n^3)$.

Proof. Clearly it suffices to prove the theorem for those n which are divisible by 18 so we will assume n/18 is integral. Let W and H be the width and height of the smallest rectangle, R, enclosing the n-permutation graph. The basic idea is to find line segments which are intersected by many disjoint edges of the permutation graph. Every intersection must be a different grid node on the line segment thus providing a lower bound on the length of the segment.

For example we first observe that $H \geq n/3$, since either there is some vertical line with n/3 terminals on it, or R can be sliced by a vertical line L into two regions containing at least n/3 terminals each. In the second case any permutation alternating between the n/3 terminals on the left and the n/3 terminals on the right, provides n/3 disjoint edges which intersect L.

A slightly more complicated example is as follows. Suppose we have a set of k disjoint paths $P_1,...,P_k$ in the grid such that each path contains 3 terminals. We will call these paths terminal paths. We will see that if there is an expanding graph of degree three on k nodes with crossing number K then some terminal path must have length at least K/k. Let G be an expanding graph of degree 3 on k nodes $v_1,...,v_k$ with crossing number K. Choose a permutation σ such that whenever v_i and v_j are adjacent in G, there is some terminal t on P_i such that $\sigma(t)$ is on P_j. Consider the cycles which realize σ. If each terminal path were contracted to a single node, the resulting graph would contain G and hence would have at least K edge crossings. Since the embeddings of the cycles are disjoint, this implies that there must be at least K intersections between terminal paths and cycles. Thus some terminal path must have length at least K/k.

The height of a terminal path is defined to be the vertical distance between the highest and lowest points on the path. Our goal is to find a set of terminal paths all with length at least $\Omega(n)$, constructed in such a way that we can add together sufficiently many of their heights to prove that $H = \Omega(n^3/W)$.

By Lemma 4.4 there is a constant c with $0 < c < 1/3$ such that for each n there is an expanding graph on $\lfloor n/37 \rfloor$ nodes with crossing number at least cn^2. We will show that $HW \geq c^2n^3$. We may assume that $W \leq cn^2$ since otherwise, as $H \geq n/3 > cn$ we clearly would have $HW \geq c^2n^3$.

Let w = 18cn. Since n/18 is integral this implies that $\lceil W/w \rceil \leq n/18$. We slice R vertically into $\lceil W/w \rceil$ slices of width at most w, and in each slice connect all of its terminals by a descending path of minimal length. We now define the terminal paths as the subpaths joining triples of terminals along the descending path. This is illustrated in Figure 6. It is easy to check that the number of terminal paths is at least $n/3 - \lceil W/w \rceil \geq n/3 - n/18 = 5n/18$. Also note that for each terminal path its height is at least its length - 2w since it is a subpath of a descending path of minimal length.

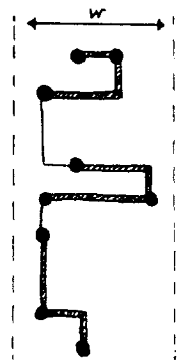

Figure 6. Constructing terminal paths

Next we discard all terminal paths that are in slices with less than nw/6W terminal paths. Since this removed at most n/6 terminal paths, at least n/9 terminal paths remain. Each remaining terminal path that is one of the $\lceil nw/12W \rceil$ longest terminal paths in its slice is labelled a long terminal path, and also removed. Now w = 18cn and $W \leq cn^2$ imply that $nw/6W \geq 3$. Thus if there were x terminal paths in a slice before removing the long terminal paths and $x > 0$, the number remaining after removing the long paths is $x - \lceil nw/12W \rceil \geq x - \lceil x/2 \rceil \geq x/3$ because $x \geq 3$. From this it is easy to check that removing the long terminal paths still leaves at least n/27 terminal paths. It follows from our previous observations that since $n/27 > n/37$, one of these terminal paths, say P, must have length at least 37cn since there is a degree 3 expanding graph on $\lceil n/37 \rceil$ nodes with crossing number at least cn^2. Now consider the long terminal paths removed from the slice containing P. Since they are all at least as long as P they each have length at least 37cn and hence height at least $37cn - 2w \geq cn$. Moreover as these terminal paths are disjoint subpaths of a descending path, the height of the slice is at least the sum of their heights, i.e. at least $(nw/12W)(cn) \geq c^2n^3/W$. Thus $HW \geq c^2n^3$ as desired. □

The proof of the $\Omega(n^3)$ lower bound for the area of n-path permutations is very similar though slightly more complicated. Because of space limitations we merely sketch how the preceding proof must be modified. The set of k node disjoint terminal paths must be replaced by k/2 input paths and k/2 output paths such that the k paths are node disjoint. Moreover the k-node expanding graph chosen must be a bipartite graph between two sets of k/2 nodes. In order to construct the input paths and output paths, we find a set of disjoint rectangular regions of bounded width, such that the union of a subset of the regions contains a fixed fraction of the inputs, and the union of the other regions contains the same fraction of the outputs. In each "input containing" region, input paths are formed exactly as the terminal paths were, by taking subpaths of a descending path of minimal length connecting all the inputs in that region. Output paths are constructed similarly in "output containing" regions. Otherwise the proof is entirely analogous.

5. PLANAR GRAPHS IN ONE LAYER

In this section we prove results about embedding planar graphs in the grid. The first is an algorithm which given an embedding of a planar graph in the plane, constructs a topologically equivalent $O(n^2)$ area (rake) grid embedding with two important properties. First, every edge makes a bounded number of turns, and secondly every exterior node has a path with a bounded number of turns to an edge of the rectangle enclosing the embedding. In fact if the exterior face is not a simple cycle and hence some exterior nodes appear more than once on the exterior face, there is such a path to an edge of the rectangle from every exterior "side" of an exterior node. Figure 7 illustrates an embedding with these properties.

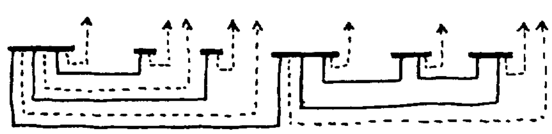

Figure 7 Paths from exterior nodes to the perimeter

In his thesis [Sh76], Shiloach gives an $O(n^2)$ area grid embedding algorithm for planar graphs but the embeddings produced have neither of these properties. Valiant [Va81] also claims an $O(n^2)$ area embedding for planar graphs but in order for his idea to work it is necessary to ensure that his embedding has the second property (which Valiant does not do). In fact, using the ideas in the proof of Theorem 4.5, it can be shown that Valiant's scheme may produce an $O(n^4)$ area embedding in the worst case.

We will use our $O(n^2)$ area embedding technique directly in proving our upper bounds for the 2AL model in the next section. However it is also implicitly needed in Theorem 3.5 since (although this observation is missing in their paper [DLT84]) Dolev, Leighton, and Trickey need an $O(n^2)$ area embedding with the second property in order to prove that every n-node planar graph of gauge at most g has an O(ng) area grid embedding.

The second result in this section shows that for any fixed placement of the nodes, a planar graph of gauge g can be embedded with this placement, so long as the specified placement places the nodes sufficiently far apart.

Although the actual $O(n^2)$ area embedding algorithm is quite simple, in order to accurately describe it and prove that it has the desired properties, we must first introduce some notation. Given a planar embedding τ of a connected planar graph G, we use the term *clockwise boundary walk produced by* τ for the ordered list of nodes visited by a complete tour of the exterior face. When the exterior face is not a simple cycle, some nodes will appear in the list more than once. Moreover the list depends on where in the exterior face the tour is started. Namely, changing the starting point results in a cyclic shift of the list. We will use the term corner to refer to two adjacent edges on the boundary of the exterior face, and refer to the corner formed by the first edge and last edge of a boundary walk as the starting corner. If the two adjacent edges are (u,v) and (v,w) we denote the corner by (u,v,w). For example 1,2,1,3,1,4,5,6,4 is a clockwise boundary walk of the graph shown in Figure 8(a), and its starting corner is (4,1,2). If the graph G is a single vertex, $\{v_1\}$, we use the convention that it has no corners and denote the starting corner of the boundary walk v_1 by ϕ (i.e. the empty corner).

Suppose $v_1,...,v_k$ is a clockwise boundary walk produced by a planar embedding τ of G with starting corner (v_k,v_1,v_2). We define $(G,\tau,(v_k,v_1,v_2))$ to be the planar supergraph of G obtained by adding k+1 new leaf nodes $x_1,...,x_{k+1}$ with x_i adjacent to v_i for $i = 1,...,k$ and x_{k+1} adjacent to v_1. We define the natural embedding of $(G,\tau,(v_k,v_1,v_2))$ as the planar embedding of $(G,\tau,(v_k,v_1,v_2))$ obtained by first embedding G with τ and adding the new edges in the exterior face so that the edge (v_i,x_i) lies in the exterior corner (v_{i-1},v_i,v_{i+1}), and (v_1,x_{k+1}) lies in the (new) exterior corner (v_k,v_1,x_1). This is illustrated in Figure 8(b). We will call the new nodes x_i the corner leaves of $(G,\tau,(v_k,v_1,v_2))$. The edges (v_1,x_{k+1}) and (v_1,x_1) can be thought of as marking the starting corner of the clockwise boundary walk. We use them to mark the corner where we wish an edge to join a node. If G is the single node, v_1, we define (G,τ,ϕ) to be the edge (v_1,x_1). This is consistent with the definition for larger G as the number of corner leaves added is always one more than the number of corners.

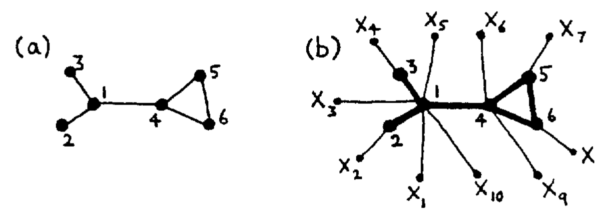

Figure 8. G and (G,τ,(4,1,2))

We say that a path in the grid is J-shaped if from the leftmost endpoint the path runs down, then right, and then up to the rightmost endpoint. Figure 9(a) shows a J-shaped path P, and two other paths P_1 and P_2 which we will describe as running along the inside and the outside of P respectively. A path is double-J-shaped if it consists of two nested J-shaped paths whose rightmost endpoints are on the same horizontal line, and a horizontal segment joining those endpoints. Figure 9(b) shows a double-J-shaped path.

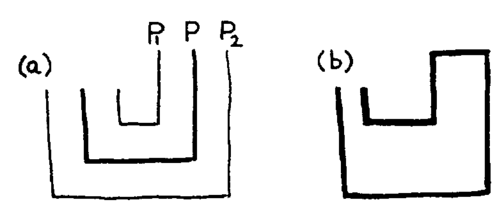

Figure 9. J-shaped and double-J-shaped paths

Recall from section 2 (see figure 1) that a rake grid embedding of a planar graph is a one layer grid embedding where nodes are represented by horizontal line segments, and edges connect into these segments from below. In a rake grid embedding, the clockwise order of edges at a node is understood to be the right to left order along the line segment representing the node.

If τ is a planar embedding of G with starting corner (v_k,v_1,v_2), we define a $(\tau,(v_k,v_1,v_2))$ embedding of G to be a rake grid embedding of $(G,\tau,(v_k,v_1,v_2))$ with the following properties:

1. It is (up to topological equivalence) the natural embedding of $(G,\tau,(v_k,v_1,v_2))$.

2. The top horizontal line contains the corner leaf nodes $x_1,\ldots,x_{k+1}$ in left to right order, but no other nodes or edges.

3. The nodes of G lie on one horizontal line with v_1 as the leftmost node.

4. Each edge (v_i,x_i) is J-shaped and (v_1,x_{k+1}) is also. Moreover the corner leaf is always the rightmost endpoint of the edge.

5. All edges of G are either J-shaped or double-J-shaped.

An example is shown below in Figure 10. If $G = \{v_1\}$, a (τ,ϕ) embedding of G is simply an embedding of the edge (v_1,x_1) as a J-shaped path so that v_1 is to the left of, and below, x_1.

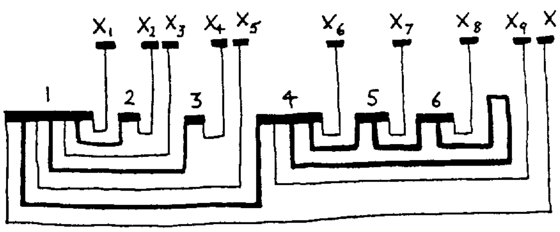

Figure 10. A $(\tau,(4,1,2))$ embedding of G

Lemma 5.1. If G has a $(\tau,(v_k,v_1,v_2))$ embedding then it has one in $O(n^2)$ area where n is the number of nodes of G.

Proof. Let m be the number of edges of G. As G is planar, m = O(n). Since each edge has a bounded number of segments the number of horizontal and vertical lines actually used by edges is O(m) = O(n). Thus unused horizontal and vertical lines can be deleted to obtain an $O(n^2)$ area $(\tau,(v_k,v_1,v_2))$ embedding of G. □

Lemma 5.2. In any $(\tau,(v_k,v_1,v_2))$ embedding of G all nodes and edges lie within the region R which is bounded by the edge from x_{k+1} to v_1, the edge from v_1 to x_1, and the segment of the top line running from x_1 to x_{k+1}.

Proof. By property (1) in clockwise order around v_1 the edge to x_{k+1} immediately precedes the edge to x_1 and hence no other edge adjacent to v_1 can leave v_1 in the complement of R. Clearly no edge can cross the boundary of R and so it suffices to show that no nodes lie in the complement of R. Suppose some node lies in the complement. Let v be the node whose distance (in the graph sense) from v_1 is minimal, and let v' be the neighbor of v on a shortest path joining v to v_1. Now if v' = v_1 then the edge from v_1 to v' leaves v_1 in the complement of R which is impossible as we observed above. On the other hand if v' is not v_1, then as the edge (v,v') cannot cross the boundary of R, v' also lies in the complement of R. However v' is closer to v_1, contradicting our choice of v. □

Theorem 5.3. For any planar embedding τ of a connected planar graph G and clockwise boundary walk $v_1,\ldots,v_k$, there is a $(\tau,(v_k,v_1,v_2))$ embedding of G.

Proof. The proof is by induction on the number of edges of G. If G has no edges then $G = \{v_1\}$, and constructing a (τ,ϕ) embedding of G is trivial. Namely we embed v_1 to the left of, and below x_1, and connect them with a J-shaped path.

Suppose G has at least one edge and that the theorem is true for all connected planar graphs with fewer edges. Let G' be the graph obtained by removing the edge (v_1,v_k) from G. There are two cases to consider depending on whether G' is connected.

If G' is disconnected it has two connected components, say G_1 and G_2, with v_1 in G_1 and v_k in G_2. For i = 1,2 let τ_i be the planar embedding of G_i obtained by restricting τ to G_i. Since the removal of (v_1,v_k) disconnects G, there exists j with $1 \le j < k$ such that $v_j = v_1$ and $v_{j+1} = v_k$. To see this note that the walk $v_1,\ldots,v_k$ must contain the edge (v_1,v_k) since otherwise the walk would connect v_1 and v_k in G' and hence G' would be connected. Also (v_1,v_k) must be traversed from v_1 to v_k since otherwise again v_1 and v_k would be connected in G'. It is easy to see that $v_1,\ldots,v_{j-1}$ is a clockwise boundary walk of G_1 produced by τ_1, and $v_{j+1},\ldots,v_{k-1}$ is a clockwise boundary walk of G_2 produced by τ_2. To obtain a $(\tau,(v_k,v_1,v_2))$ embedding of G, we take the $(\tau_1,(v_{j-1},v_1,v_2))$ embedding of G_1 and the $(\tau_2,(v_{k-1},v_j,v_{j+2}))$ embedding of G_2, and lay them side by side with G_1 on the left so that all the nodes of G are on the same horizontal line. If necessary one of the embeddings is stretched at the top so that the corner leaves of both graphs are embedded on the top line. Label (in left to right order) the j corner leaves of the left graph $x_1,\ldots,x_j$ and the k-j corner leaves of the right graph $x_{j+1},\ldots,x_k$. Figure 11(a) illustrates the placement of these graphs. Note that by Lemma 5.2 all edges and nodes lie within the two shaded regions shown in Figure 11(a). The final corner node is added at the top right corner and the two remaining edges, (v_1,v_k) and (v_1, x_{k+1}) are routed as shown in Figure 11(b).

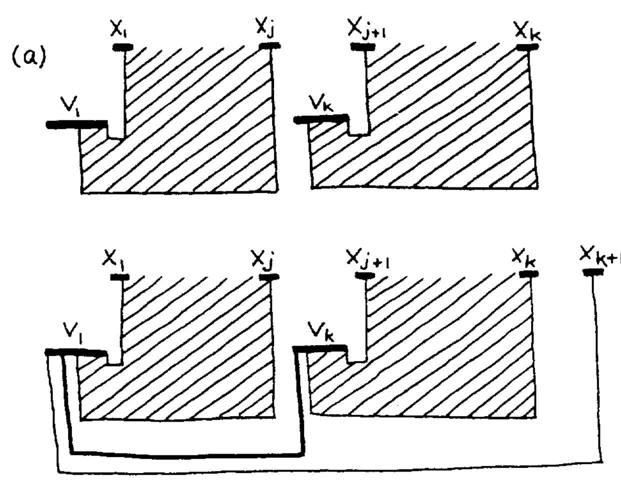

Figure 11. The $(\tau,(v_k,v_1,v_2))$ embedding when G' disconnected

It is not hard to see that the graph we have embedded is

$(G,\tau,(v_k,v_1,v_2))$, and the embedding is the natural embedding. It is also easy to check that the remainder of the properties are satisfied, so that indeed this is a $(\tau,(v_k,v_1,v_2))$ embedding of G.

Now suppose G' is connected, and let τ' be the restriction of τ to G'. Then τ' produces a clockwise boundary walk $v_1,\ldots,v_k,v_{k+1},\ldots,v_t$ where (in counter-clockwise order) $v_1,v_k,v_{k+1},\ldots,v_t$ is the interior face of G containing the edge (v_1,v_k). To obtain the $(\tau,(v_k,v_1,v_2))$ embedding of G, we take the $(\tau',(v_t,v_1,v_2))$ embedding of G', delete the corner leaves $x_{k+1},\ldots,x_t$ and the edges incident to them, change the name of x_{t+1} to x_{k+1}, and route the edge from v_1 to v_k along the inside of edge (v_1,x_{k+1}), left until just before x_k, then along the outside of the edge (x_k,v_k) until reaching v_k. Finally, the corner leaf nodes $x_1,x_2,\ldots,x_{k+1}$ are raised up one line so that property 2 is satisfied. This is illustrated in Figure 12. Again it is easy to see that this is the natural embedding of $(G,\tau,(v_k,v_1,v_2))$, and that all the other properties are satisfied. □

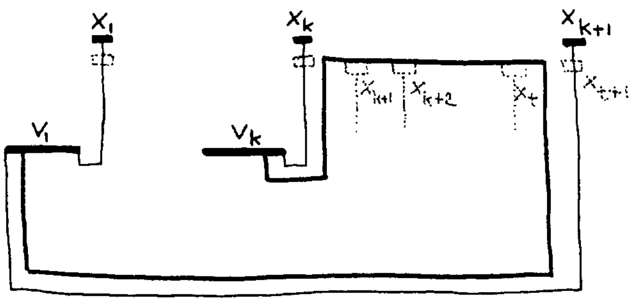

Figure 12. The $(\tau,(v_k,v_1,v_2))$ embedding when G' connected

The proof we have just given obviously yields an algorithm for embedding a planar graph in $O(n^2)$ area such that each edge makes at most 6 turns, and such that each exterior corner of the graph has a path with at most 2 turns to the boundary of the enclosing rectangle. We will call this embedding the J embedding of the graph. Given an embedding τ and a starting corner (u,v,w) we define the $(\tau,(u,v,w))$ ordering of the nodes of G to be the left to right order given by the J embedding. (If G is disconnected choose some arbitrary left to right ordering of its connected components.)

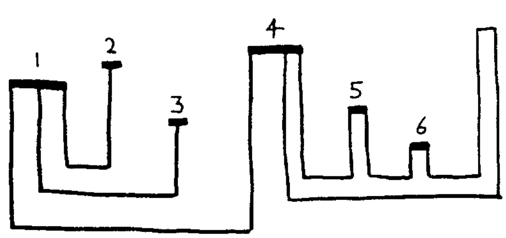

Figure 13. A stretched J embedding

Remark 5.4. Let τ be a planar embedding of an n-node planar graph G and let (u,v,w) be a starting corner. Then for any grid placement of the nodes of G such that the left to right ordering of the nodes is the $(\tau,(u,v,w))$ ordering, after "stretching" the grid placement by inserting O(n) new horizontal and vertical lines, G can be embedded with the stretched placement.

Proof. Since the J embedding only uses O(n) vertical lines and O(n) horizontal lines, it suffices to show how the J embedding can be vertically stretched so that the top to bottom ordering of the nodes agrees with that of the grid placement. To the J embedding, add n new horizontal lines immediately above the horizontal line containing the nodes of G. Now each node is slid up onto one of the new horizontal lines so that each node lies on a different horizontal line and so that the top to bottom ordering agrees with the grid placement. An example is shown in Figure 13. □

Lemma 5.5. Let G be an n-node planar graph of bounded degree with gauge g. Then for any grid placement of G with at most one node position of G on each horizontal line, after stretching the placement by inserting $O(gn^2)$ new horizontal lines and $O(g^2n)$ new vertical lines, G can be embedded with the stretched placement.

Proof. First insert $O(n^2)$ new horizontal lines so that each node position is at least n lines below the top of the rectangle and each pair of node positions is separated by a vertical gap of at least n lines. Next, add n new vertical lines on both sides of the rectangle. Call the top line the gluing line, and this rectangle the extended rectangle. Let τ be an embedding of G such that each node in G has a path of length at most g to the exterior face. Next take the J embedding of G with respect to τ and any starting corner, turn it upside down, and place it on top of the gluing line. Since every node has a path of length at most g to an exterior node it is easy to see that in the inverted J embedding of G, every node has a path to the gluing line with O(g) turns. Moreover each path only crosses O(g) edges in the inverted J embedding. By inserting O(gn) new vertical and horizontal lines these paths to the gluing line can be taken to be disjoint (though they will still generally intersect edges of the inverted J embedding). Thinking of the endpoints of these paths on the gluing line as the n inputs and the node positions in the extended rectangle as the n outputs, Remark 4.2 observes that the extended rectangle is an n-path permutation graph. Thus using the ideas of 4.1 it is possible to extend the paths from the nodes in the inverted J embedding to the node positions in the extended rectangle so that all the paths are disjoint and each path connects the node in the inverted J embedding to its correct position in the extended rectangle. Figure 14 illustrates our progress so far.

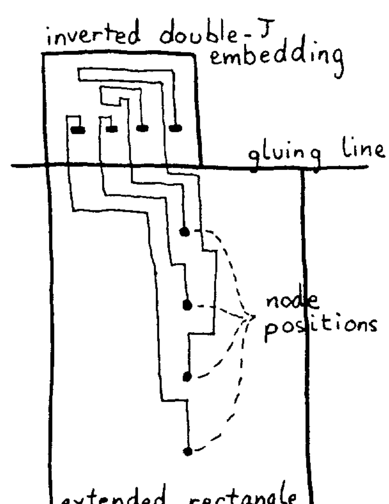

Figure 14. Preliminary stages of the embedding

Now for each path, the edges of the inverted J embedding which cross it are pulled around its node position endpoint in the extended rectangle so that they no longer cross it. This is shown in Figure 15. Of course in doing this new horizontal and vertical lines must be inserted. Within the inverted J embedding $O(g^2n)$ new vertical lines and horizontal lines are needed as there are n paths, each path makes $O(g)$ turns and crosses $O(g)$ edges. Within the extended rectangle $O(gn)$ new vertical lines and $O(gn^2)$ new horizontal lines are needed. Since $g < n$, at most $O(gn^2)$ new horizontal lines and $O(g^2n)$ new vertical lines are needed altogether.

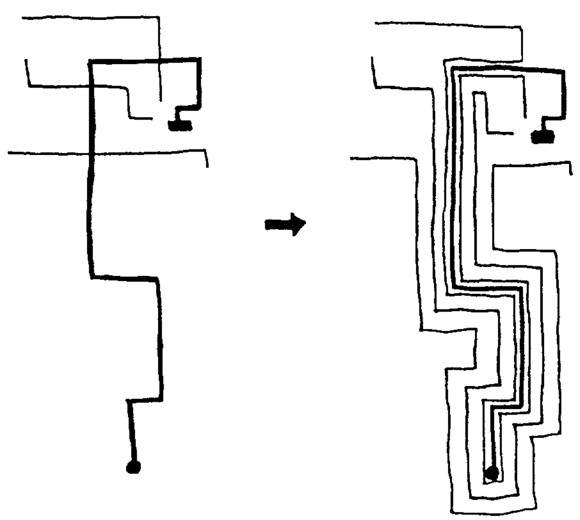

Figure 15. Removing edge crossings

Finally each node of G in the inverted J embedding and the edges leaving it are pulled along its path to its position in the extended rectangle. Since there are $O(n)$ edges in G, in the inverted J embedding $O(gn)$ new vertical and horizontal lines are needed and in the extended rectangle $O(n)$ new vertical lines and $O(n^2)$ new horizontal lines are needed. Putting all this together we see that $O(gn^2)$ new horizontal lines and $O(g^2n)$ vertical lines were inserted. □

Remark 5.6. Let G be an n-node planar graph. Then for any grid placement of G with at most one node position of G on each horizontal line, after stretching the placement by inserting $O(n^3)$ new horizontal lines and $O(n^2)$ new vertical lines, G can be embedded with the stretched placement.

Proof. The embedding is identical to that described in the proof of 5.5 except that in constructing the paths joining the nodes in the inverted J embedding to their positions in the extended rectangle, the initial part of the path lying inside the inverted J embedding is taken just to be the straight vertical line joining the node to the gluing line. This line has at most $O(n)$ crossings with edges in the inverted J embedding since G has $O(n)$ edges and each edge makes a bounded number of turns. Pulling these edges around the paths requires $O(n^2)$ new vertical lines within the inverted J embedding, and in the extended rectangle $O(n^3)$ new horizontal lines and $O(n^2)$ new vertical lines are needed. The requirements for pulling the edges connected to a node along its path are as above except that only $O(n)$ new vertical lines are needed in the inverted J embedding. □

6. EMBEDDING GRAPHS IN k ACTIVE LAYERS

Theorem 6.1. Every n-node graph of thickness at most 2 can be rake embedded in two active layers in $O(n^2)$ area using no contact cuts.

Proof. Let G be an n-node graph which is the union of two planar n-node graphs, G_1 and G_2. For $i = 1,2$ choose an embedding τ_i of G_i in the plane and a starting corner (u_i,v_i,w_i). Let σ_i be the $(\tau_i,(u_i,v_i,w_i))$ ordering of the nodes of G. It suffices to find a placement of the nodes of G in the n x n grid such that the left to right ordering is σ_1 and the top to bottom ordering is σ_2, since by 5.4 with the insertion of $O(n)$ new vertical and horizontal lines we can embed G_1 in the first layer. Since inserting the new lines does not change the top to bottom ordering (nor left to right for that matter), we can now apply 5.4 again (rotated 90 degrees) to embed G_2 in the second layer with the insertion of another $O(n)$ new vertical and horizontal lines. Let r_i and c_i be the position of the i-th node of G in σ_2 and σ_1 respectively. It is easy to see that placing the i-th node in row r_i and column c_i achieves the desired left to right and top to bottom orderings. □

Theorem 6.2. There are n-node graphs of degree at most 3 and thickness at most 2 such that for any fixed k every k active layer embedding requires $\Omega(n^2)$ area, regardless of the number of contact cuts.

Proof. It suffices to find such a graph of degree at most 3, since we will see in 6.4 and 6.5 that every graph of degree at most 3 has thickness at most 2. Let G be an n-node expanding graph of degree at most 3, and assume we have an embedding of G in k active layers. By increasing the area by at most a factor of k^2, we may assume that no pair of horizontal [vertical] segments on different layers occupy the same horizontal [vertical] line. By projecting each layer onto the first layer, we obtain a planar embedding of G such that any pair of distinct edges only intersect at grid nodes, and at most 2 edges intersect at any grid node. Now since G has $\Omega(n^2)$ crossing number by Lemma 4.4, there must be $\Omega(n^2)$ grid nodes where two edges cross each other, which shows that the embedding has $\Omega(n^2)$ area. □

The reason we can embed graphs of thickness 2 in $O(n^2)$ area in 2 active layers without contact cuts, is that since the grid is two dimensional we can achieve two different orderings of the nodes simultaneously. However for graphs of thickness k in k active layers when $k \geq 3$, the situation seems to be more difficult. If G is an n-node graph which is the union of k n-node planar graphs, the best we can do at the moment is to position each node on a different horizontal line and then by Remark 5.6, with the insertion of at most $O(n^3)$ horizontal lines and $O(n^2)$ vertical lines, we can embed one of the planar graphs on each layer. This yields the following theorem.

Theorem 6.3. Every graph of thickness at most k has an $O(n^5)$ area rake embedding in k active layers without contact cuts.

For graphs of bounded degree we can do somewhat better by partitioning the edges into cycle covers and using the fact that there is an $O(n^3)$ area grid which is an n-cycle permutation graph. More precisely, a cycle cover of a graph is a set of node disjoint cycles such that every node of the graph is in some cycle. We assume that the graph may have loops and multiple edges and that cycles of length one or two are permitted. The following theorem is due to J. Peterson [Pe91].

Theorem 6.4. If G is a 2k-regular graph then its edges can be partitioned into k edge disjoint cycle covers.

Simple arguments yield the following lemma.

Lemma 6.5. If G is an n-node graph of degree at most 2k, there is an n-node 2k-regular graph (which may have multiple edges and loops) which contains G as a subgraph.

Theorem 6.6. Any n-node graph of degree at most 2k can be embedded in the k active layer model in $O(n^3)$ area using no contact cuts.

Proof. By 6.5 we may assume the graph is 2k-regular, and hence by 6.4 the edges can be partitioned into k cycle covers. We position the n nodes in the $4n^2 \times 4n$ grid exactly as the terminals are positioned in the proof of 4.3. Now using 4.3 we can embed each cycle cover in a different layer. □

Remark 6.7. It is easy to see how these ideas can be extended to give an $O(nb^2)$ embedding for graphs of bandwidth b. Thus graphs with small separators can be embedded much more efficiently.

Since planar graphs can be embedded in $O(n^2)$ area on one layer, one expects that there should be better embeddings without contact cuts for planar graphs in several active layers. The remainder of this section outlines an $O(n^{1.6})$ area embedding for n-node planar graphs of bounded degree in two active layers without contact cuts.

The basic idea is as follows. The edges of the graph are partitioned into two sets, say E_1 and E_2, which will be embedded on the first and second layers respectively. The edges in E_1 are embedded using the Dolev-Leighton-Trickey algorithm which embeds any n-node planar graph of gauge g in a square with O(ng) area (see 3.1). Since some nodes are endpoints of edges in both E_1 and E_2, the embedding of the edges in E_1 on the first layer places constraints on how the edges of E_2 are to be embedded in the second layer. Fortunately the edges in E_2 form a graph of constant gauge and we can use Lemma 5.5 to embed them without too disastrous an effect on the area.

Theorem 6.8. Every n-node planar graph of bounded degree has an $O(n^{1.6})$ area embedding in two active layers without contact cuts.

Proof. Let G be an n-node planar graph of bounded degree. Applying 3.4 with $g = n^{.6}$ we can partition the edges of G into E_1 and E_2 so that the induced subgraphs, G_1 and G_2, have gauges at most $n^{.6}$ and 4 respectively, and E_2 has $O(n^{.4})$ edges. By 3.1 G_1 can be embedded on the first layer in a square with side S where $S = O(n^{.8})$. Let N be the set of nodes of G incident with edges in E_2. Since $|N| = O(n^{.4})$ we may assume also that the embedding of G_1 places each node of N on a different horizontal line. Now by 5.5 we can embed the edges in E_2 on the second layer with the insertion of at most $O(n^{.8})$ new horizontal lines and $O(n^{.4})$ new horizontal lines, thus achieving the promised $O(n^{1.6})$ area. □

7. OPEN PROBLEMS

As the results in this paper show, many questions about these models remain to be answered. Here we mention a few open problems which we consider important.

7.1. Close the gap between the $\Omega(n^2)$ lower bound and $O(n^5)$ upper bound for the area required to embed a graph of thickness k in k active layers when $k \geq 3$.

7.2. Leighton's $\Omega(n \log n)$ bound for the area required to embed planar graphs in the grid, when crossings of horizontal and vertical segments of edges are allowed [Leig81], can be used to prove an $\Omega(n \log n)$ lower bound on the area required to embed planar graphs in k active layers even with contact cuts, for any fixed k. Can planar graphs be embedded without contact cuts in less than $O(n^{1.6})$ area when more than two active layers are allowed? Can the $\Omega(n \log n)$ lower bound on the area needed to embed planar graphs in two layers be improved when contact cuts are not allowed?

7.3. Give upper bounds on the area needed to embed planar graphs in k active layers which depend on the gauge of the planar graph.

7.4. How large must a planar (not necessarily embedded in the grid) permutation network be? Without the grid embedding assumption the current lower bound is the obvious one of $\Omega(n^2)$ which follows directly from the planar separator theorem. On the other hand all known examples of planar permutation networks have size $\Omega(n^3)$.

Acknowledgements

We are grateful to Noga Alon, Ashok Chandra, Ruth Cohen, Danny Dolev, Rao Kosaraju, Tom Leighton, Gary Miller, and Michael Rabin, for comments and suggestions. 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.

REFERENCES

[BP77] N. K. Bose and K. A. Prabhu, "Thickness of Graphs with Degree Constrained Vertices", IEEE Trans. on Circuits and Systems, Vol. CAS-24, No. 4, April 1977, pp. 184-190.

[C78] F. R. K. Chung, "On Concentrators, Superconcentrators, Generalizers and Nonblocking Networks", Bell Sys. Tech. J., 58 (1978) 1765-1777.

[CS78] M. Cutler and Y. Shiloach, "Permutation Layout", Networks, Vol. 8, 1978, pp. 253-278.

[DLT84] D. Dolev, F. T. Leighton, and H. Trickey, "Planar Embeddings of Planar Graphs", Advances in Computing Research, Vol. 2, pp. 147-161, JAI Press, 1984.

[KM82] S. Rao Kosaraju and G. Meyer, Personal Communication.

[KKF79] E. S. Kuh, T. Kashiwabara, and T. Fujisawa, "On Optimum Single Row-Routing", IEEE Trans. on Circuits and Systems, Vol. CAS-26, No. 6, June 1979, pp. 361-368.

[Ka80] Y. Kajitani, "On Via Hole Minimization of Routings on a Two-Layer Board", Tech. Report, Dept. of Elec. Engineering, Tokyo Institute of Technology, Japan, 1980.

[Lei81] F. T. Leighton, "Layouts for the Shuffle-Exchange Graph and Lower Bound Techniques for VLSI", Ph. D. thesis, Dept. of Math., MIT, 1981.

[Leig81] F. T. Leighton, "New Lower Bound Techniques for VLSI", Proceedings 22nd Ann. Symp. Found. of Comp. Science, 1981, pp. 1-12.

[Lei82] F. T. Leighton, "A Layout Strategy Which is Provably Good", Proc. 14th Annual ACM Symp. on the Theory of Computing, 1982, pp. 85-92.

[Lei85] F. T. Leighton, Personel Communication.

[Lei80] C. E. Leiserson, "Area Efficient Graph Layouts (for VLSI)", Proc. 21st Ann. IEEE Symp. on Foundations of Computer Science, 1980, pp. 270-281.

[LT79] R. J. Lipton and R. E. Tarjan, "A Separator Theorem for Planar Graphs", SIAM J. Appl. Math., 36 (1979) 177-189.

[LT80] R. J. Lipton and R. E. Tarjan, "Applications of a Planar Separator Theorem", SIAM J. Comp., 9 (1980) 615-627.

[M] J. Makowsky, "On the Complexity of Certain Decision Procedures for Propositional Calculus and Random Graphs", undated preprint, Free University, Berlin and Hebrew University, Jerusalem.

[MC80] C. A. Mead and L. A. Conway, "Introduction to VLSI Systems", Addison Wesley, Reading Mass., 1980.

[Mi84] G. Miller, "Finding Small Simple Cycle Separators for 2-Connected Planar Graphs", Proc. 16th Annual ACM Symp. on Theory of Computing, 1984, pp. 369-376.

[Pe91] J. Peterson, "Die Theorie der regulaeren Graphen", Acta Mathematica Vol. 15, 1891, pp.193-220.

[Ri82] R. Rivest, "The Placement and Interconnect System", Proc. 19th Design Automation Conference, 1982, pp. 475-481.

[Sh76] Y. Shiloach, "Linear and Planar Arrangement of Graphs", Ph. D. Dissertation, Weizmann Institute, Rehovot, Isreal, 1976.

[Sh80] I. Shirakawa, "Letter to the Editor : Some Comments on Permutation Layout", Networks, Vol. 10, 1980, pp. 179 - 182.

[So74] H.C. So, "Some Theoretical Results on the Routing of Multilayer Printed Wiring Boards", 1974 IEEE Int. Symp. on Circuits and Systems, pp. 296-303. [Th80] C. D. Thompson, "A Complexity Theory for VLSI", Ph. D. Dissertation, Carnegie-Mellon University, 1980.

[TKS76] 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. of Circuits and Systems, Vol. CAS-23, No. 12, 1979, pp. 768-778.

[TK82] S. Tsukiyama and E. S. Kuh, "Double-Row Planar routing and Permutation Layout", Networks, Vol. 12, 1982, pp. 287-316.

[Va81] L. G. Valiant, "Universality Considerations in VLSI", IEEE Trans. on Computers, Vol. C-30, No. 2, Feb. 1981, pp. 135-140.