Asymptotic Component Densities in Programmable Gate Arrays Realizing all Circuits of a Given Size*

T. Berger[1,2] A. Hekstra[1] A. Orlitsky[2]

[1]Cornell University [2]AT&T Bell Laboratories

Abstract

A Programmable Gate Array (PGA) is modeled as a square grid. Some grid nodes are processing nodes containing electrical elements. The rest are switching nodes capable of connecting wires incident on them. Two possible types of switching nodes are considered. In vertex connectivity each switching node can connect only one pair of wires. In edge connectivity each switching node can simultaneously connect two pairs of wires. The PGA must be capable of implementing any graph of size at most k and degree at most four. We prove tight bounds on the highest achievable density of processing nodes.

In edge connectivity the highest achievable density is $\Theta(1/k)$. In vertex connectivity the highest achievable density is $\Theta(1/k^2)$. If the grid is augmented by the diagonal edges then the highest achievable density is $\Theta(1/k)$ even with vertex connectivity. These extend known results for embedding graphs in grids.

Small graphs of degree one are further examined. For $k = 2$ and $k = 3$ the highest density of processing nodes equals the highest density of parked cars in a square parking lot where each car can exit. Both densities are $2/3$. For $k = 4$ the highest density is $1/2$.

*Appeared in Algorithmica 9(2):101-127, February 1993. This copy was printed on November 20, 1995.

1 Introduction

A *Programmable Gate Array (PGA)* consists of many components placed on a wafer and of wires capable of connecting them. By selecting a subset of the components and connecting the appropriate wires, various circuits can be implemented. A PGA thereby combines the versatility of a printed circuit board with the speed of a single-wafer VLSI chip.

Ideally, a PGA would have a wire between any two components. Then every circuit could be implemented by connecting the appropriate wire ends to the components. However, the number of wires in such a PGA would grow as the square of the number of components.

Therefore, a compromise usually is adopted whereby the PGA resembles a graph we call the *PGA graph*. The wires are the edges of the PGA graph; the points where they can be connected are the nodes. Some of the nodes contain the electrical elements; these are the *processing nodes*. Each processing node can connect the electrical element it contains to any of the wires incident upon it, but cannot directly connect two wires[1]. The other nodes are *switching nodes*. They can connect pairs of wires incident upon them. We distinguish between two assumptions concerning the connection capabilities of the switching nodes:

Vertex Connectivity — Each switching node can connect only one pair of wires.

Edge Connectivity — Each switching node acts as a crossbar and can simultaneously connect several pairs of wires.

These names were chosen in part because in vertex connectivity the switching nodes can be regarded as connecting the two wires to the node (vertex), while in edge connectivity the switching nodes can be regarded as connecting pairs of wires (edges) directly; however, the main reason for adopting these names will become clear in the next section. In circuits implemented by PGA's with vertex connectivity, the wires correspond to vertex-disjoint paths in the PGA graph whereas in circuits implemented by PGA's with edge connectivity, the wires correspond to edge-disjoint paths. Edge connectivity clearly requires a more advanced technology. However, one of our main conclusions is that edge connectivity allows for a significant increase in the number of components the PGA can contain.

Remark: Previous works on embedding graphs in grids considered one-layer graph embeddings and two-layer graph embeddings. One-layer graph embedding is equivalent to vertex connectivity. Two-layer graph embedding is similar to edge connectivity except that

[1]Another model where processing nodes can connect wires is discussed in Section 7.

wires can run in two layers, hence the two are equivalent if we ignore constant factors pf the number of wires. In current PGA's, cf. [Fr 88] and [EGRREM 88], the only freedom afforded the user is to determine switching-node connections; hence we use vertex and edge connectivities.

Circuits, too, can be modeled as graphs, called *circuit graphs*. The vertices correspond to the electrical elements of the circuit and the edges represent the connections. For a PGA to *implement* a circuit, the processing and switching nodes are reconfigured by the user. Two processing nodes are connected by a chain of PGA wires if and only if the corresponding electrical elements are connected in the circuit. We assume that each electrical element of the circuit corresponds to a unique component of the PGA. This is the case, for example, if each type of electrical element appears in the PGA exactly once[2].

Although any graph can serve as a PGA graph, we shall consider only the standard *two-dimensional square grid*. The grid is popular, among other reasons, for its regularity, its high connectivity, and its simple layout. Figure 1 illustrates a PGA. The processing nodes are numbered and the switching nodes are drawn empty.

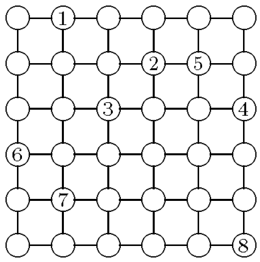

Figure 1: A 6 by 6 grid PGA with 8 processing nodes.

Example 1 Assume that the PGA graph depicted in Figure 1 has to embed the circuit graph $1-2$, $3-4$, $5-6$, where a dash represents a required connection. With vertex connectivity we can connect the wires as shown by the thick lines in Figure 2(a). With edge connectivity we can either use the same paths or the simpler paths shown in Figure 2(b). Note that the path form 3 to 4 crosses the path from 5 to 6 at a switching node. This is permissible in edge-connectivity PGA's. If the circuit contained the additional wire $7-8$, it

[2]This assumption constitutes a major departure from most real PGA's. For further discussion, see Section 7.

could not have been implemented with vertex-connectivity but could have been implemented with edge connectivity. □

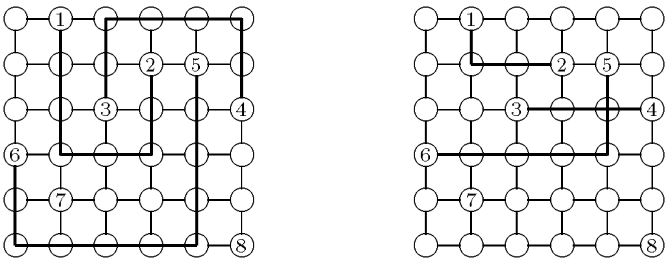

Figure 2: Embedding a circuit graph in a PGA graph.

It is desirable to have many processing nodes in the PGA. This, however, limits the circuits that can be implemented, both because there are fewer switching nodes left for connecting the wires and because, if the processing nodes are too close together, not all circuits using these nodes can be implemented[3]. It is this tradeoff between density and connectivity that we shall investigate. For a given k, we want to determine the largest proportion of grid nodes that can be used as processing nodes subject to the requirement that every k-element circuit using these nodes can be implemented.

We state the problem formally in the next section. We define $\nu_{d,k}$ to be the largest proportion of grid nodes that can serve as processing nodes in a vertex-connectivity PGA if every circuit of degree d and size k whose elements are among the processing nodes can be implemented by the PGA. We define $\epsilon_{d,k}$ similarly for edge-connectivity PGA's.

We assume that each edge of the PGA graph represents one actual wire. Since every grid node has degree at most four, the PGA cannot implement circuits of degree larger than four[4]. Therefore $\epsilon_{d,k} = 0$ for $d \geq 5$.

Recall that when a PGA implements a circuit, each circuit wire is mapped into a sequence of wires, or path, in the PGA. In a vertex-connectivity PGA these paths cannot intersect

[3]We shall see later that this is the more dominant of the two reasons. Hence, for large PGA's there is a diminishing difference between PGA's with processing nodes that can connect wires and PGA's where the processing nodes cannot connect wires. See Section 7.

[4]In real PGA's, each grid edge represents several wires. Circuits of higher degree can therefore be implemented. This is discussed in Section 7.

at any node, so the crossing number of the circuit graph cannot exceed the crossing number of the PGA graph. Grids are planar graphs; hence, vertex-connectivity grid PGA's can implement only planar circuits. Since there are nonplanar graphs of degree three, $\nu_{d,k} = 0$ for $d \geq 3$. Define $\nu_{d,k}^{\text{planar}}$ as $\nu_{d,k}$ except that only planar graphs need be embedded. Clearly, $\nu_{1,k}^{\text{planar}} = \nu_{1,k}$ and $\nu_{2,k}^{\text{planar}} = \nu_{2,k}$.

The simplest type of circuit that involves k elements consists of pairs of connected elements. Such a circuit has degree 1. First, we consider embedding degree-1 graphs of small size. Degree-1 graphs with 2 or 3 vertices have only one edge, hence the same PGA's can implement both. Also, for one embedded edge, there is no advantage to edge connectivity. Therefore $\nu_{1,2}$, $\nu_{1,3}$, $\epsilon_{1,2}$, and $\epsilon_{1,3}$ coincide. In Section 3 we draw on an analogy to parking lots to prove that

$$\nu_{1,2} = \nu_{1,3} = \epsilon_{1,2} = \epsilon_{1,3} = \frac{2}{3} \ .$$

We argue that the largest density of parked cars in a square lot is precisely $\nu_{1,2}$. The standard parking lot arrangement consists of two parked rows, then one empty row, followed by two more parked rows, and so on. We show that this arrangement, which clearly achieves a $\frac{2}{3}$ density of parked cars, is asymptotically optimal.

In Section 4 we consider embedding degree-1 graphs of size 4 in vertex-connectivity PGA's. We show that

$$\nu_{1,4} = \frac{1}{2} \ .$$

We then turn to embedding graphs of general fixed size. In Section 5 we consider embedding size-k circuit graphs in edge-connectivity PGA's. Clearly $\epsilon_{0,k} = 1$, and we remarked earlier that $\epsilon_{d,k} = 0$ for $d \geq 5$. We show that for all k

$$\frac{1}{8(k+1)} \leq \epsilon_{4,k} \leq \epsilon_{3,k} \leq \epsilon_{2,k} \leq \epsilon_{1,k} \leq \frac{32}{k} \ .$$

That is, for $1 \leq d \leq 4$

$$\epsilon_{d,k} \in \Theta\left(\frac{1}{k}\right) \ .$$

Hence, regardless of whether the PGA must implement just the simplest kind of circuits (pairs of processors) or the most complicated ones (arbitrary degree-4 circuits), the highest possible density of processing nodes is asymptotically of the same order.

In grid graphs each vertex has degree ≤ 4. Hence with edge connectivity, any switching node can simultaneously connect at most two pairs of wires. With vertex connectivity each switching node can connect one pair. Although this difference might appear inconsequential, we show in Section 6 that it accounts for a large difference in achievable density. Again, $\nu_{0,k} = 1$ and $\nu_{d,k} = 0$ for $d \geq 5$, but for all k,

$$\frac{c_1}{k^2} \leq \nu_{4,k}^{\text{planar}} \leq \nu_{3,k}^{\text{planar}} \leq \nu_{2,k} \leq \nu_{1,k} \leq \frac{c_2}{k^2} \ .$$

where c_1 and c_2 are constants. That is, for $1 \leq d \leq 4$,

$$\nu_{d,k}^{\text{planar}} \in \Theta\left(\frac{1}{k^2}\right) \ .$$

This shows that as k, the size of the circuits guaranteed implementable, increases, the component densities achievable in vertex and edge connectivity PGA's decrease at different rates. Hence, the capability to connect more than one pair of wires at a switching node allows for a significant increase in density.

The results of Sections 5 and 6 are related to those of [CS 78], [AKLLW 85], and [AKS 89]. These papers consider the minimal size of a grid that can embed planar graphs of size n with prescribed vertex locations. Ignoring slight model differences, their results show that in the model equivalent to vertex connectivity the grid has to be of size $\Theta(n^3)$ while in the model equivalent to edge connectivity the required size is $\Theta(n^2)$.

These results can be derived from those described in Sections 5 and 6. Let n denote both the number of processing nodes and the size of the embedded graphs[5]. For edge connectivity, we show that $\frac{n}{\text{grid size}} = \Theta\left(\frac{1}{n}\right)$, therefore grid size $= \Theta(n^2)$. For vertex connectivity, $\frac{n}{\text{grid size}} = \Theta\left(\frac{1}{n^2}\right)$, therefore grid size $= \Theta(n^3)$. Thus, the results of Sections 5 and 6 can be viewed as generalizing known results to the embedding of graphs of size k where k is not necessarily the number of processing nodes.

Most of the paper is devoted to the model discussed in this introduction. The PGA graph is a two dimensional grid. Each edge represents one wire. The electrical elements are located in the processing nodes. A processing node can only connect the element it contains to a wire; it cannot connect any pair of wires. In Section 7 we discuss alternative models. We show that the asymptotic results remain valid even if the processors are placed in the squares rather than at the nodes and even if each can connect to all vertices around it. They are unchanged even if the processing nodes can connect pairs of wires.

We show that adding the diagonal edges to the grid, which only doubles the number of wires (see Figure 18), increases the asymptotic component density in vertex-connectivity PGA's to $\Theta\left(\frac{1}{k}\right)$, that which is achievable with edge-connectivity grid PGA's (with or without

[5]The proofs have to be slightly modified to account for increasing graph size.

the diagonal edges). As we shall see later on, this occurs partly because the resulting PGA graph is no longer planar.

2 Definitions

A *graph* is an ordered pair (V, E). V is a set, called the *vertex set*; its elements are *vertices* or *nodes*. E is a collection of *edges*: unordered pairs of vertices. An edge consisting of two identical vertices is a *self edge*, and an edge appearing more than once in E is a *multiple edge*. The following standard definitions assume that $G = (V, E)$ is a graph. Strictly speaking, all definitions should be qualified by the underlying graph (e.g., *vertices of* G) but, for brevity, G is omitted throughout.

If u and v are vertices and $\{u, v\}$ is an edge, u and v are *adjacent*, *neighbors*, or *connected by* $\{u, v\}$, and the edge $\{u, v\}$ *connects* u and v. A *path* is a sequence $v_0, \ldots, v_l$ of vertices such that v_i is adjacent to v_{i+1} for $i = 0, \ldots, l-1$. The vertices v_0 and v_l are the *end vertices* of the path and the vertices $v_1, \ldots, v_{l-1}$ are the *interior vertices*. The edges $\{v_i, v_{i+1}\}$ for $i = 0, \ldots, l-1$ are the *edges of the path*. Let u and v be vertices of G. A *path between* u *and* v, or a *path connecting* u *and* v, is a path whose end vertices are u and v. u and v are *connected* if there is a path between them. Note that a single vertex is a path; hence a vertex is always connected to itself even if it is not contained in a self loop. Let $S \subseteq V$ be a set of vertices. A *path in* S is a path whose interior vertices are in S. The end vertices may or may not be in S.

Two paths $u_0, \ldots, u_l$ and $v_0, \ldots, v_m$ are *vertex disjoint* if they do not share any vertices: $u_i \neq v_j$ for all $0 \leq i \leq l$ and all $0 \leq j \leq m$. The paths are *edge disjoint* if they do not share any edge. A set of paths is *vertex disjoint (edge disjoint)* if every pair of two paths in the set is vertex disjoint (edge disjoint).

Let $G' = (V', E')$ be a graph with $V' \subseteq V$ and let $S \subseteq V$. A *vertex-disjoint embedding of* G' *in* S is a mapping that, with each edge $\{v_1, v_2\} \in E'$, associates a path in S connecting v_1 and v_2 such that the resulting set of paths is vertex disjoint. If such an embedding exists, G' is *vertex-disjoint embeddable in* S. An *edge-disjoint embedding of* G' in S, and being *edge-disjoint embeddable in* S, are defined similarly except that the paths have to be edge disjoint.

Recalling the Programmable Gate Array model described in the introduction, think of G as the PGA graph with E as the set of wires and V as the set of nodes, and think of S as the set of switching nodes. The graph G' represents a circuit we want the PGA to implement. Each vertex of G' is an electrical element corresponding to a unique component, or vertex, of V. Note that the vertices of G' are fixed. They correspond to specific components in the PGA. The only freedom in the embedding is in the choice of the interior vertices of the paths. In vertex-connectivity PGA's, the paths cannot cross at a vertex. Hence the circuit G' can be implemented by the PGA if and only if G' is vertex-disjoint embeddable in S. In the edge-connectivity model, G' can be implemented by the PGA if and only if G' is edge-disjoint embeddable in S.

The *size* $|S|$ of a set S is the number of its elements. The *size* of a graph is the size of its vertex set. A graph is *size-k* if its size is at most k. The *degree* of a vertex is the number of edges containing it. A self edge counts twice and multiple-edge counts add. The *degree* of a graph is the largest degree of its vertices. A graph is *degree-d* if its degree is at most d. The complement of a subset S of vertices is $V - S \stackrel{\text{def}}{=} \{v : v \in V \text{ and } v \notin S\}$.

Let d and k be positive integers. A subset P of V is *(d,k)-vertex-disjoint embedding* if every degree-d, size-k graph with vertices in P is vertex-disjoint embeddable in $V-P$. (Note that it is P that we call(d,k)-vertex-disjoint embedding although the graphs are embeddable in $V-P$.) The *(d,k)-vertex-disjoint density* of a graph G is the size of the largest (d,k)-vertex-disjoint embedding set in G, normalized by the size of G:

$$\nu_{d,k}(G) \stackrel{\text{def}}{=} \frac{\max\{|P| : P \text{ is } (d,k)\text{-vertex-disjoint embedding }\}}{|V|} \ .$$

In the PGA model, G is the PGA graph, P represents the processing nodes, and $V-P$ is the set of switching nodes, assumed vertex connecting. P is (d,k)-vertex-disjoint embedding if and only if every degree-d, size-k circuit whose electrical elements correspond to components of P can be implemented by the PGA. The (d,k)-vertex-disjoint density of G is the largest possible proportion of processing nodes in a vertex-connecting PGA based on G if this PGA can implement every degree-d, size-k circuit.

We define a *(d,k)-edge-disjoint embedding* set similarly, and let the *(d,k)-edge-disjoint density* of a graph be:

$$\epsilon_{d,k}(G) \stackrel{\text{def}}{=} \frac{\max\{|P| : P \text{ is } (d,k)\text{-edge-disjoint embedding }\}}{|V|} \ .$$

For all d and k and all graphs G,

$$\nu_{d,k}(G) \leq \epsilon_{d,k}(G) \ ,$$

and if $d' \geq d''$ and $k' \geq k''$ then

$$\nu_{d',k'}(G) \leq \nu_{d'',k''}(G) \ .$$

Because of the application motivating this paper, we are not interested in the full generality of (d,k)-vertex-disjoint or (d,k)-edge-disjoint embeddings. Rather, we restrict our attention to the *2-dimensional square grid graphs*, which model most existing PGA's[6]. For every $m, n \geq 1$, the *2-dimensional m by n grid* denoted by $G_{m,n} \stackrel{\text{def}}{=} (V_{m,n}, E_{m,n})$ is the graph whose vertex set is $V_{m,n} \stackrel{\text{def}}{=} \{0,\ldots,m-1\} \times \{0,\ldots,n-1\}$ and whose edges are the *horizontal* edges: $\{\{(i,j-1),(i,j)\} : i = 0,\ldots,n-1,\ j = 1,\ldots,n-1\}$ and the *vertical* edges: $\{\{(i-1,j),(i,j)\} : i = 1,\ldots,n-1,\ j = 0,\ldots,n-1\}$; cf. Figure 3. Note that the first coordinate corresponds to the *vertical* axis and that coordinate values increase from top to bottom and from left to right. These conventions are adopted throughout the paper. We call $G_{n,n}$ a *square grid* and denote $G_{n,n}$, $V_{n,n}$, and $E_{n,n}$ by G_n, V_n, and E_n respectively.

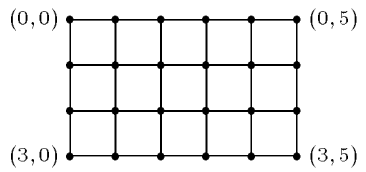

Figure 3: $G_{4,6}$.

A *planar representation* of a graph G is a drawing in the plane consisting of a dot corresponding to each vertex of G and a simple curve in the plane connecting two dots if and only if the corresponding vertices are connected by an edge in G. At most two lines are allowed to intersect at any point in the plane. Every graph has a planar representation. The *crossing number* of a graph G is the minimal number of line crossings in a planar representation of G. A graph with crossing number 0 is *planar*.

If G' is vertex-disjoint or edge-disjoint embeddable in a subset of V, then the degree of every vertex $v \in V'$ in G' is not larger than its degree in G. Therefore in grid graphs $\epsilon_{d,k}(G_n)$ is nonzero only for $d \leq 4$. In practice, each edge of G_n often represents several actual wires, in which case circuits of degree larger than four can be implemented[7].

[6] For several variations, see Section 7.

[7] See Section 7.

Furthermore, with vertex connectivity when a graph G' is embedded in another graph G, each edge of G' is mapped into a path in G and distinct paths do not intersect at any node. A planar representation of G can therefore be used to derive a planar representation of G' with at most as many line crossings. Hence, in vertex connectivity, a graph can be embedded only in a graph with at least as large a crossing number. Grid graphs are planar and therefore can only embed planar graphs. Since there is a non-planar graph of six vertices and degree three, $\nu_{d,k}(G_n)$ is zero for $d \geq 3$ and $k \geq 6$.

We are interested in $\nu_{d,k}(G_n)$ for large grids. Therefore we define the (d,k)-*vertex-disjoint density*:

$$\nu_{d,k} \stackrel{\text{def}}{=} \limsup_{n\to\infty} \nu_{d,k}(G_n) \ .$$

Analogously, we define the (d,k)-*edge-disjoint density*:

$$\epsilon_{d,k} \stackrel{\text{def}}{=} \limsup_{n\to\infty} \epsilon_{d,k}(G_n) \ .$$

$\nu_{d,k}$ has the interpretation of the largest proportion of grid vertices that can serve as processing nodes in a vertex-connectivity PGA if every degree-d, size-k circuit can be implemented by the PGA. $\epsilon_{d,k}$ has the same interpretation for edge-connectivity PGA's.

A description and a brief interpretation of the results obtained in the paper were given in the last part of the introduction.

3 Embedding Degree-1 Graphs of Size 2 or 3

Consider embedding degree-1 graphs with 2 or 3 vertices. Such graphs have at most one edge. Therefore, a set is $(1,2)$-vertex-disjoint embedding if and only if it is $(1,3)$-vertex-disjoint embedding (similarly for edge-disjoint embeddings). Furthermore, there is no advantage to a path that intersects itself; hence, vertex-disjoint and edge-disjoint connectivity coincide and

$$\nu_{1,2} = \nu_{1,3} = \epsilon_{1,2} = \epsilon_{1,3} \ .$$

We determine $\nu_{1,2}$. To obtain a lower bound, we draw on an analogy. Consider a parking lot consisting of an n by n array of square spaces. One of the spaces on the perimeter is marked EXIT. Each space either contains a parked car or is vacant. Cars can move from a space to any of the four adjacent spaces provided it is vacant. The parking lot can accommodate n^2 cars but then only one car can move to the EXIT space: the car that is there already. It is desirable to allow each car to "EXIT" without moving any other car. How many cars can be parked?

The parking lot problem is closely related to $(1,2)$-vertex-disjoint embeddability. An n by n parking lot with p parked cars can be used to derive a $(1,2)$-vertex-disjoint embedding set P of size p in G_n. Consider the natural correspondence between G_n and an n by n parking lot in which each node of G_n represents the corresponding parking space. Let P consist of the nodes corresponding to spaces containing parked cars, as illustrated in Figure 4. (Note that the arrangement shown is not valid because the car marked "1" cannot EXIT.) Since every car can EXIT, there is a path in $V_n - P$ from every node in P to the EXIT node (the node corresponding to the EXIT space). Hence, there is a path in $V_n - P$ between any two nodes in P. The set P of processing nodes is therefore $(1,2)$-vertex-disjoint embedding and of size p. The converse, though not used (we prove the upper bound directly), is also correct. If a set $P \subseteq V_n$ is $(1,2)$-vertex-disjoint embedding, eliminate from P a node closest to the EXIT node; the parking spaces corresponding to the remaining nodes can be used to park cars. Therefore, as n increases, the proportion of cars that can be parked in an n by n parking lot approaches $\nu_{1,2}$.

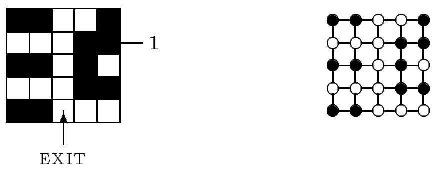

Figure 4: An arrangement of parked cars and the corresponding PGA.

The standard parking-lot arrangement is depicted in Figure 5. Two pairs of parked columns are separated by one column of vacant spaces. For n divisible by 3, the total number of parked cars is $\frac{2}{3}n(n-1)+2$. Hence,

$$\nu_{1,2}(G_n) \geq \frac{2}{3} - \frac{2}{3n} + \frac{2}{n^2} ,$$

implying

$$\nu_{1,2} \geq \frac{2}{3} .$$

In Theorem 1 we shall show that

$$\nu_{1,2} \leq \frac{2}{3} ,$$

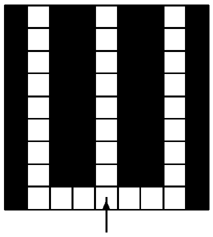

Figure 5: Standard arrangement of a square parking lot.

thereby proving that the standard arrangement asymptotically achieves optimal density. First, we need some preliminaries.

Let $S \subseteq V$ be a set of vertices. Two vertices (not necessarily in S) are *connected in S* if there is a path in S that connects them. As a property, being connected in S is reflexive and symmetric but, in general, not transitive. For example, in the graph $u - v - w$, where edges are denoted by $-$, the vertices u and v are connected in $S \stackrel{\text{def}}{=} \{u, w\}$, as are v and w. However, the only path connecting u and w contains the interior vertex v, so u is not connected to w in S. However, when restricted to elements of S, the property of being connected in S is also transitive, hence an equivalence relation. Therefore, a set S is partitioned into equivalence classes of *S-connected components*. A set S of vertices is *self connected* if every two of its elements are connected in S. The following properties hold for any set S:

P1. If two vertices are connected in S, they are connected in one of the S-connected components.

P2. Each of the S-connected components is self connected.

P3. S is self connected if and only if it has only one S-connected component.

P4. S is self connected if and only if $G|_S$, the *restriction* of G to S, is connected. ($G|_S$ is the graph whose vertex set is S and in which $u, v \in S$ are connected by an edge if and only if $\{u, v\}$ is an edge in E.)

Let $S \subseteq V$. The *vicinity of S* is the set of vertices outside S that are adjacent to vertices of S:

$$\mathcal{V}(S) \stackrel{\text{def}}{=} \{v \in V - S : v \text{ is adjacent to some } s \text{ in } S\} .$$

Lemma 1 If S is self connected in a two-dimensional grid, then

$$|\mathcal{V}(S)| \leq 2|S| + 2 .$$

Proof: By induction on $|S|$. If $|S| = 1$ then S has at most 4 neighbors. If S is self connected and has more than one element, there is an element $s \in S$ such that $S - \{s\}$ is self connected (S can be chosen as a leaf of a spanning tree for $G|_S$). By the induction hypothesis, $|\mathcal{V}(S - \{s\})| \leq 2(|S| - 1) + 2 = 2|S|$.

Let $T \stackrel{\text{def}}{=} (S - \{s\}) \bigcap \mathcal{V}(\{s\})$ be the set of vertices in $S - \{s\}$ that are neighbors of s. T is not empty because S is self connected and has at least 2 elements. Clearly,

$$\mathcal{V}(S) = \Big(\mathcal{V}(S - \{s\}) \bigcup \mathcal{V}(\{s\})\Big) - \Big(T \bigcup \{s\}\Big) .$$

Therefore,

$$\begin{aligned} |\mathcal{V}(S)| &\leq |\mathcal{V}(S - \{s\})| + |\mathcal{V}(\{s\})| - |\mathcal{V}(S - \{s\}) \bigcap \mathcal{V}(\{s\})| - |T| - 1 \\ &\leq 2|S| + 4 - 0 - 1 - 1 \\ &= 2|S| + 2. \quad \square \end{aligned}$$

Theorem 1

$$\nu_{1,2} = \frac{2}{3} .$$

Proof: We argued earlier that $\nu_{1,2} \geq \frac{2}{3}$. Now we show that $\nu_{1,2} \leq \frac{2}{3}$. Let S and P be disjoint sets of vertices in G_n such that P is connected in S. Pick a vertex $p \in P$ and let $S_1, S_2, \ldots, S_I$ be the S-connected components with elements adjacent to p (that is, $p \in \mathcal{V}(S_i)$ for $1 \leq i \leq I$). These components are disjoint and p has at most 4 neighbors; hence $0 \leq I \leq 4$. Since P is connected in S, any element of P is S-connected to p. From property 3 above, any element of P is S_i-connected to p for some $1 \leq i \leq I$. But a vertex v is S_i connected to p if and only if $v \in \mathcal{V}(S_i)$ or $v \in \mathcal{V}(p)$. Therefore, $P \subseteq \mathcal{V}(p) \bigcup \left(\bigcup_{i=1}^{I} \mathcal{V}(S_i)\right)$. By the previous lemma,

$$\begin{aligned} |P| &\leq (4 - I) + \sum_{i=1}^{I} |\mathcal{V}(S_i)| - (I - 1) \\ &\leq (4 - I) + \sum_{i=1}^{I} 2|S_i| + 2I - (I - 1) \\ &\leq 2|S| + 5 . \end{aligned}$$

But $|P| + |S| \leq n^2$. Combined, the last two inequalities yield: $|P| \leq 2(n^2 - |P|) + 5$, or

$$|P| \leq \frac{2}{3}n^2 + \frac{5}{3} .$$

Therefore, $\nu_{1,2} = \lim_{n\to\infty} \nu_{1,2}(G_n) \leq \lim_{n\to\infty} \frac{2}{3} + \frac{5}{3n^2} = \frac{2}{3}$. $\square$

4 Embedding Degree-1 Graphs of Size 4

Unlike graphs of size two or three, embedding graphs of size 4 may be different with vertex and edge connectivity. We consider only vertex connectivity.

The subset of V_n consisting of the center $n-2$ elements in every other column (see Figure 6) is easily seen to be $(1,4)$-vertex-disjoint embedding, so $\nu_{1,4} \geq \frac{1}{2}$. We prove a matching lower bound thereby showing that

$$\nu_{1,4} = \frac{1}{2} \; .$$

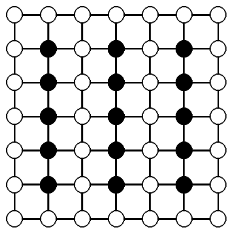

Figure 6: A $(1,4)$-vertex-disjoint embedding set of size 15 in G_7.

Let $\mathcal{Z}$ denote the set of integers. Define the *south-eastern neighborhood of a point* $s = (i,j) \in \mathcal{Z}^2$ to be

$$SE(s) \stackrel{\text{def}}{=} \{(i+1,j),(i,j+1),(i+1,j+1)\} \; ,$$

the set of grid points, 'south,' 'east,' and 'south east' of s. Define the *south-eastern neighborhood of a set* $S \subseteq \mathcal{Z}^2$ to be

$$SE(S) \stackrel{\text{def}}{=} \bigcup_{s \in S} SE(s) - S \; ,$$

the set of points in $\mathcal{Z}^2 - S$ that are 'south,' 'east,' or 'south east' of some point in S. Figure 7 illustrates a set and its south-eastern vicinity. The elements of the set are denoted by dots and those of the south-eastern neighborhood by $\times$'s. Lemma 5, proved below, shows that if P is $(1,4)$-vertex-disjoint embedding, then $|SE(P)| \geq |P|$. This enables us to prove

Theorem 2

$$\nu_{1,4} \leq \frac{1}{2} \; .$$

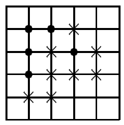

Figure 7: A set and its south-eastern neighborhood

Proof: Let P be $(1,4)$-vertex-disjoint embedding in G_n. By definition, P and $SE(P)$ are disjoint and their union is contained in $\{0,\ldots,n\} \times \{0,\ldots,n\}$. Therefore,

$$|P| + |SE(P)| \leq (n+1)^2$$

From Lemma 5, $|P| \leq |SE(P)|$, so

$$\frac{|P|}{n^2} \leq \frac{1}{2} + \frac{1}{n} + \frac{1}{2n^2}\ .$$

This holds for all n and all $(1,4)$-vertex-disjoint embedding sets in G_n. Letting n tend to infinity, we obtain the theorem. $\square$

To prove that $|SE(P)| \geq |P|$, we first show that $(1,4)$-vertex-disjoint embedding sets are comprised of simple subsets. Let $d(s_1,s_2) \stackrel{\text{def}}{=} ((i_1-i_2)^2 + (j_1-j_2)^2)^{1/2}$ denote the *Euclidean distance* between two points $s_1 = (i_1,j_1)$ and $s_2 = (i_2,j_2)$ in $\mathcal{Z}^2$. A mapping $\phi : \mathcal{Z}^2 \to \mathcal{Z}^2$ is *distance-preserving* or an *isometry* if $d(s_1,s_2) = d(\phi(s_1),\phi(s_2))$ for all $s_1,s_2 \in \mathcal{Z}^2$. Two subsets S and T of $\mathcal{Z}^2$ are *isometric* if there is an isometry from S onto T. Since any isometry is $1-1$, isometric sets have the same cardinality. (Also, two isometric sets are related via a sequence of rotations, translations, and reflections.)

The restriction $G_n|_P$ of the square grid G_n to a set $P \subseteq V_n$ was defined in Section 3. Let $G_1(P),\ldots,G_k(P)$ be the connected components of $G_n|P$ (see Figure 8), and for $i = 1,\ldots,k$ let P_i be the vertices of $G_i(P)$.

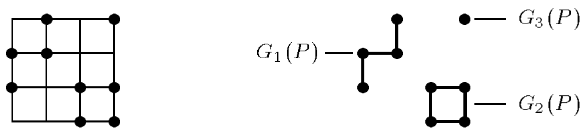

Figure 8: A set P in G_4 and its connected-components.

Define $\mathcal{Z}^+ \stackrel{\text{def}}{=} \{0,1,2,\ldots\}$ and $\mathcal{Z}^- \stackrel{\text{def}}{=} \{0,-1,-2,\ldots\}$. The following lemma can be proved by inspection.

Lemma 2 If P is $(1,4)$-vertex-disjoint embedding in G_n then each P_i is isometric to a connected subset of one of the following sets:

1. $S_L \overset{\mathrm{def}}{=} \mathcal{Z}^- \times \{0\} \bigcup \{0\} \times \mathcal{Z}^+$
2. $S_T \overset{\mathrm{def}}{=} \{-1\} \times \mathcal{Z} \bigcup \{(0,0)\}$
3. $S_Z \overset{\mathrm{def}}{=} \{0\} \times \mathcal{Z}^+ \bigcup \{-1\} \times \mathcal{Z}^-$.

Subsets of S_L, S_T, and S_Z are illustrated in Figure 9 which also suggests the reason why these three subscripts where selected. □

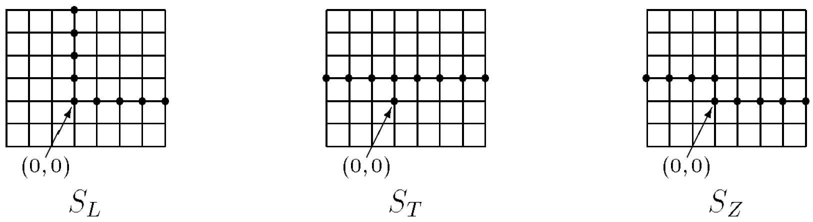

Figure 9: The different types of $(1,4)$-vertex-disjoint embedding connected components.

Lemma 3 If C is a $(1,4)$-vertex-disjoint embedding connected component, then

$$|SE(C)| = |C| + 2 \ .$$

Proof: By induction on the size of C. There are fourteen types of connected components: the four rotations of each of S_L, S_T, and S_Z, the 'horizontal' connected component, and the 'vertical' one. Induction is carried out separately for each. Figure 10 shows four of the inductions bases. □

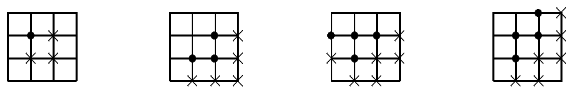

Figure 10: A few induction bases showing that $|SE(C)| = |C| + 2$.

Let $P \subseteq V_n$ be $(1,4)$-vertex-disjoint embedding and let $P_1, \ldots, P_k$ be the vertex sets of P's connected components as defined earlier. For each $i \in \{1, \ldots, k\}$ define S_i, T_i, and U_i as follows:

$$
\begin{aligned}
S_i &\stackrel{\text{def}}{=} \{x \in SE(P_i) : x \notin SE(P_j) \text{ and } x \notin P_j \text{ for all } j \neq i\}, \\
T_i &\stackrel{\text{def}}{=} \{x \in SE(P_i) : x \in SE(P_j) \text{ for some } j \neq i\}, \\
U_i &\stackrel{\text{def}}{=} \{x \in SE(P_i) : x \in P_j \text{ for some } j \neq i\}.
\end{aligned}
$$

For every $i \in \{1, \ldots, k\}$, the sets S_i, T_i, and U_i are disjoint and their union is $SE(P_i)$. First, we bound the number of points in T_i and U_i.

Lemma 4 Let $P \subseteq V_n$ be $(1,4)$-vertex-disjoint embedding. For $i \in \{1, \ldots, k\}$ define P_i, S_i, T_i, and U_i as above. Then, for all $i \in \{1, \ldots, k\}$,

$$\frac{|T_i|}{2} + |U_i| \leq 2 \ .$$

Proof: All fourteen cases mentioned in the proof of Lemma 3 need to be examined. We illustrate four of them here.

Let P_i be the connected component depicted in Figure 11(a). Only the point u_1 may belong to U_i and only t_1 and t_2 may belong to T_i. Hence $|T_i| \leq 2$ and $|U_i| \leq 1$, so $\frac{|T_i|}{2} + |U_i| \leq 2$. Note that if $t_3 \in T_i$, then $v \in P$, and the connections $v - 2$, $1 - 3$ cannot be made simultaneously; similarly, $t_4 \notin T_i$.

Substituting u_2 for t_4 at the end of the last sentence, the argument above also shows that $\frac{|T_i|}{2} + |U_i| \leq 2$ for the connected component P_i depicted in Figure 11(b).

For the connected component P_i depicted in Figure 11(c), only the point u_1 may belong to U_i and only t_1, t_2, and t_3 may belong to T_i. If $\frac{|T_i|}{2} + |U_i| > 2$, then $u_1 \in U_i$ and all of t_1, t_2, t_3 are in T_i. In particular, u_1 and v are both in P. But then, regardless of the other points in P, the connections $u_1 - 1$ and $v - 2$ cannot be made simultaneously.

For the connected component P_i depicted in Figure 11(d), only the points u_1 and u_2 may belong to U_i, and only t_1 and t_2 may belong to T_i. If $\frac{|T_i|}{2} + |U_i| > 2$, then u_1 and u_2 are both in U_i and at least one of t_1 and t_2, say t_1, is in T_i. But then u_1, u_2, and v are in P and, regardless of the other points in P, the connections $u_1 - 1$ and $v - 2$ cannot be made simultaneously. $\Box$

Lemma 5 Let P be $(1,4)$-vertex-disjoint embedding in G_n. Then

$$|SE(P)| \geq |P| \ .$$

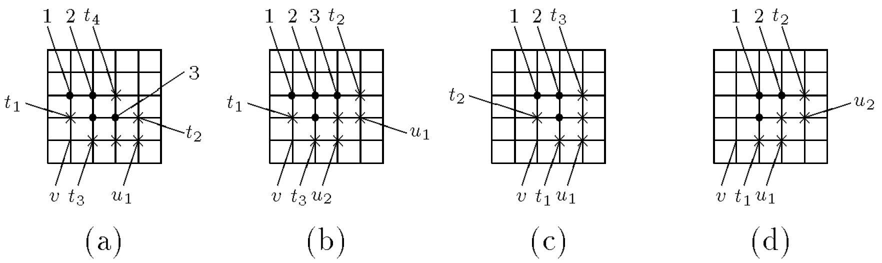

Figure 11: Four connected components satisfying $\frac{|T_i|}{2} + |U_i| \leq 2$.

Proof: Let $P_1, \ldots, P_k$ be the vertex sets of the connected components of P and define S_i, T_i, and U_i for $i \in \{1, \ldots, k\}$ as above. An element of V_n belongs to $SE(P)$ if and only if it belongs to T_i or U_i for some $i \in \{1, \ldots, k\}$. Hence,

$$SE(P) = \bigcup_{i=1}^{k}(S_i \bigcup T_i) = (\bigcup_{i=1}^{k} S_i) \bigcup (\bigcup_{i=1}^{k} T_i) \ .$$

By definition, all the S_i's are disjoint and each element of $\bigcup_{i=1}^{k} T_i$ belongs to T_i for exactly two indices i. Hence, using Lemmas 3 and 4

$$\begin{aligned} |SE(P)| &= \sum_{i=1}^{k} |S_i| + \frac{\sum_{i=1}^{k} |T_i|}{2} \\ &= \sum_{i=1}^{k} (|S_i| + |T_i| + |U_i|) - (\frac{|T_i|}{2} + |U_i|) \\ &\geq \sum_{i=1}^{k} (|P_i| + 2) - 2 \\ &= |P| \ . \end{aligned}$$

□

5 Edge-Disjoint Embeddings of Size-k Graphs

In this section we assume that the switching nodes are edge connecting. Each can simultaneously connect two wire pairs. We show that for all $k \geq 1$,

$$\frac{1}{8(k+1)} \leq \epsilon_{4,k} \leq \epsilon_{3,k} \leq \epsilon_{2,k} \leq \epsilon_{1,k} \leq \frac{32}{k} \ .$$

That is, for all $1 \leq d \leq 4$,

$$\epsilon_{d,k} \in \Theta\left(\frac{1}{k}\right) \ .$$

As mentioned in the introduction, $\epsilon_{d,k} = 0$ for $d \geq 5$ and $\epsilon_{0,k} = 1$. We begin by proving the upper bound.

Theorem 3 For all $k \geq 1$,

$$\epsilon_{1,k} \leq \frac{32}{k} \ .$$

Proof: Let $n > \frac{k}{4}\max\{\frac{1}{\sqrt{2}}, \frac{\sqrt{k}}{8}\}$ and let $P \subset V_n$ be $(1,k)$-edge-disjoint embedding in G_n. We assume that

$$\frac{|P|}{n^2} > \frac{32}{k}$$

and derive a contradiction. As illustrated in Figure 12, partition V_n into $\frac{2|P|}{k}$ subsets[8], each a subgrid of dimensions at most $n\sqrt{\frac{k}{2|P|}}$ by $n\sqrt{\frac{k}{2|P|}}$.

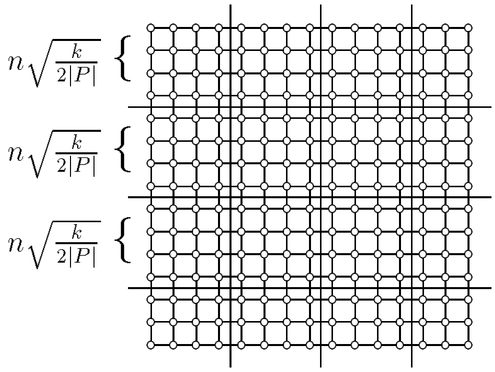

Figure 12: Partition of G_n into subgrids.

At least one of these $\frac{2|P|}{k}$ subgrids contains $\frac{k}{2}$ elements of P. Also, by the assumptions just made,

$$|P| > \frac{32}{k}n^2 > \frac{32}{k}\left(\frac{k}{4\sqrt{2}}\right)^2 = k$$

and

$$|P|^2 > \left(\frac{32}{k}n^2\right)^2 > n^2\left(\frac{k\sqrt{k}}{32}\right)^2\left(\frac{32}{k}\right)^2 = n^2k \ .$$

Hence,

$$n^2\frac{k}{2|P|} < \frac{|P|}{2} < |P| - \frac{k}{2} \ ,$$

so there must be at least $\frac{k}{2}$ elements of P outside any subgrid.

Consider a subgrid with at least $\frac{k}{2}$ elements of P. Since P is $(1,k)$-edge-disjoint embedding, there must be $\frac{k}{2}$ edge-disjoint paths, each connecting an element of P in the subgrid

[8] For simplicity we disregard integer roundoffs (e.g. $\left\lceil \frac{2|P|}{k} \right\rceil$ subsets). Since the derived bound is asymptotic (n is arbitrarily large), it is not affected by these approximations.

and an element of P outside the subgrid. Therefore, there must be at least $\frac{k}{2}$ edges connecting the subgrid to the rest of the grid. But the number of such edges is bounded above by the *circumference* of the subgrid which is at most $4n\sqrt{\frac{k}{2|P|}}$. Hence,

$$\frac{k}{2} \leq \text{circumference} \leq 4n\sqrt{\frac{k}{2|P|}}$$

which contradicts our initial assumption. □

Next, we prove a matching lower bound by applying a combination of standard embedding and 'snaking' techniques (e.g., [CS 78], [AKLLW 85]) to our problem.

Theorem 4 For all $k \geq 1$,

$$\epsilon_{4,k} \geq \frac{1}{8(k+1)} .$$

Proof: First we show that the set $P \stackrel{\text{def}}{=} \{(1, 4i+1) : i = 0, \ldots, n-1\}$ is $(4,k)$-edge-disjoint embedding in the grid $G_{2k+2,4n}$. Then we show how to "map" this rectangular grid into the square grid $G_{2(2k+2)+\lceil\sqrt{8n(k+1)}\rceil}$. This implies

$$\epsilon_{4,k} \geq \limsup_{n\to\infty} \frac{n}{\left(2(2k+1) + \left\lceil\sqrt{8n(k+1)}\right\rceil\right)^2} = \frac{1}{8(k+1)} .$$

Figure 13 shows $G_{2k+2,4n}$, the set P for $k = n = 4$, and an embedding of a sample graph. When embedding a degree-4, size-k graph G', we sequentially embed each edge (in any order). Suppose that an edge connecting the i'th and the j'th nodes of G' has not yet been embedded. Since G' is of degree ≤ 4, at least one of the four lines leaving vertex $(1, 4i+1)$ and one of the four lines leaving vertex $(1, 4j+1)$ are not used. We use these lines to get from the vertices to the first row that has not been used and connect the paths in that row. The row will not be used again.

The process is repeated until all edges have been embedded in $G_{2k+2,4n}$. Since G' is of size k and of degree at most 4 it has at most $2k$ edges so $2k+2$ rows suffice. Therefore, G' can be embedded in $G_{2k+2,4n}$.

To map $G_{2k+2,4n}$ into $G_{2(2k+2)+\left\lceil\sqrt{8n(k+1)}\right\rceil}$, we 'fold' it as shown in Figure 14 into a rectangular grid at most $\left\lceil\sqrt{8n(k+1)}\right\rceil$ by $\left\lceil\sqrt{8n(k+1)}\right\rceil$ in size. The added strips of width $2k+1$ ensure that the original paths can be rerouted in the folded version. The resulting grid has side at most $2(2k+1) + \left\lceil\sqrt{8n(k+1)}\right\rceil$. □

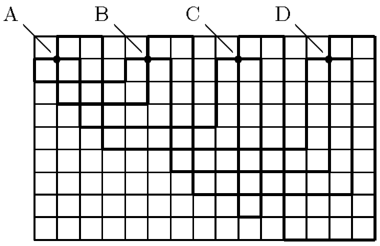

Figure 13: Embedding the graph A-B, A-B, A-C, A-D, B-D, B-D, C-C, C-D, in $G_{10,16}$.

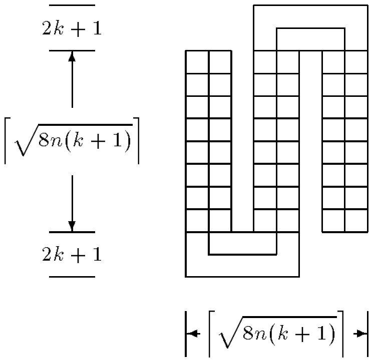

Figure 14: Mapping a rectangular grid into a square grid.

6 Vertex-Disjoint Embeddings of Size-k Graphs

In Section 5 we considered edge-connectivity PGA's. Each switching node could connect two wire pairs. We showed that the (d, k)-edge-disjoint density is

$$\epsilon_{d,k} \in \Theta\left(\frac{1}{k}\right) .$$

Now, we assume that the switching nodes of the PGA are vertex-connecting, that is, each can connect only one pair of wires. We show that this seemingly minor distinction accounts for a decided difference in achievable densities to the effect that for all k,

$$\frac{c_1}{k^2} \leq \nu_{4,k}^{\text{planar}} \leq \nu_{3,k}^{\text{planar}} \leq \nu_{2,k} \leq \nu_{1,k} \leq \frac{c_2}{k^2} .$$

where c_1 and c_2 are constants. That is, for $1 \leq d \leq 4$

$$\nu_{d,k}^{\text{planar}} \in \Theta\left(\frac{1}{k^2}\right) .$$

Thus as the size of implementable circuits increases, the achievable densities in vertex and edge connectivity PGA's decrease at different rates. Switching nodes capable of simultaneously connecting two wire pairs significantly increase the achievable density.

We do not describe the lower bound here. [AKS 89] provided an algorithm for embedding a planar graph of size n in a grid of area $O(n^3)$ with prescribed vertex locations. This algorithm can be modified easily to embed any planar graph of size k in a rectangular vertex-connectivity PGA of area $O(nk^2)$ with n processing nodes. The rectangular grid can then be mapped as in the last section into a square grid of the same area. Therefore,

$$\nu_{d,k} = \Omega\left(\frac{n}{nk^2}\right) = \Omega\left(\frac{1}{k^2}\right) .$$

The matching upper bound,

$$\nu_{d,k} \leq \frac{c_2}{k^2} ,$$

uses a lemma of [Le 84] for proving lower bounds on crossing numbers of graphs (stated here as Lemma 8) and a technique similar to one used by [AKLLW 85]. The proof, given in Theorem 5, is based on *chorded cycle graphs* and on two lemmas. These are described first.

For even k, a *chorded cycle of size k* (or a *chorded k-cycle*) is a graph whose vertex set is $\{1,\ldots,k\}$ and whose edges are the *cycle edges*: $\{\{k,1\},\{1,2\},\ldots,\{k-1,k\}\}$, and the *chord edges*: arbitrary non-self edges. Each vertex belongs to exactly one chord. An example of a chorded 8-cycle is shown in Figure 15.

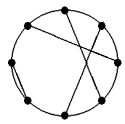

Figure 15: A chorded 8-cycle

The crossing number of a graph G was defined in Section 2 as the minimal number of line crossings in a planar representation of G. A graph with crossing number zero is planar. The graph consisting of just the cycle edges of a chorded k-cycle is planar as is the graph consisting of the chords only. The chorded k-cycle itself is not necessarily planar. For example, the crossing number of the chorded 8-cycle in Figure 15 is one.

Since a chorded k-cycle has $\frac{k}{2}$ chords, its crossing number is at most $\binom{k/2}{2} \leq \frac{k^2}{8}$. The first lemma states that there is a chorded k-cycle with a crossing number asymptotically that high. Because the proof is lengthy, it is postponed until after Theorem 5.

Lemma 6 For all even $k \geq 6$, there exists a chorded k-cycle with crossing number $\geq k^2/10^8$.

The second lemma is the grid-graph equivalent of a well known geometrical result. If $|P|$ points are placed in the unit square, there is a path of length at most $c \cdot \sqrt{|P|}$ going through all the points. The exact constant multiplying $\sqrt{|P|}$ has been studied in several papers including [Fe 55] and [Ka 89], but is not of interest here. We prove a weak version of the lemma.

Let G be a graph. A path $v_0, \ldots, v_l$ in G is *simple* if no vertex appears twice in it: $v_i \neq v_j$ for all $0 \leq i < j \leq l$. The *length* of the path is l, the number of edges in the path.

Lemma 7 Let $G_n = (V_n, E_n)$ be the square grid and let $P \subseteq V_n$. For every $1 \leq k \leq |P|$ there is a subset $p_0, \ldots, p_k$ of P and a simple path of length $\leq \frac{k}{|P|} n(4\sqrt{|P|}+2)$ in G_n connecting $p_0, \ldots, p_k$.

Proof: We prove that there is a simple path of length $\leq n(2\sqrt{|P|}+1)$ connecting all points in P. The lemma then follows from averaging arguments. Partition V_n into $\sqrt{|P|}$ 'vertical stripes' $S_1, \ldots, S_{\sqrt{|P|}}$ where $S_i \stackrel{\text{def}}{=} \{0, \ldots, n-1\} \times \{(i-1)\frac{n}{\sqrt{|P|}}, \ldots, i\frac{n}{\sqrt{|P|}}-1\}$. Let $a_i \stackrel{\text{def}}{=} |P \bigcap S_i|$ be the number of elements of P in the ith stripe and let $p^i_j = (x^i_j, y^i_j)$ be the jth element of $P \bigcap S_i$ where the elements are ordered by increasing x value for even i and decreasing x values for odd i (elements with the same x value are ordered by increasing y value). For even i, let $p^i_0 \stackrel{\text{def}}{=} (0, (i-1)\frac{n}{\sqrt{|P|}})$ and let $p^i_{a_i+1} \stackrel{\text{def}}{=} (n-1, i\frac{n}{\sqrt{|P|}}-1)$. For odd i reverse the definition of p^i_0 and $p^i_{a_i+1}$. For $i = 1, \ldots, \sqrt{|P|}$, the path connects p^i_j to p^i_{j+1} for $j = 0, \ldots, a_i$ by first going 'horizontally' then 'vertically' if necessary. The path then connects $p^i_{a_i+1}$ to p^{i+1}_0. Therefore, for each i, the length of the path in stripe S_i is at most

$$\begin{aligned}
\sum_{j=0}^{a_i} \left(|x^i_j - x^i_{j+1}| + |y^i_j - y^i_{j+1}| \right) &\leq \left| \sum_{j=0}^{a_i} (x^i_j - x^i_{j+1}) \right| + \sum_{j=0}^{a_i} \frac{n}{\sqrt{|P|}} \\
&= |x^i_{a_i+1} - x^i_0| + (a_i+1)\frac{n}{\sqrt{|P|}} \\
&= n + (a_i+1)\frac{n}{\sqrt{|P|}} \ .
\end{aligned}$$

The total length of the path is at most

$$\sum_{i=1}^{\sqrt{|P|}} (n + (a_i+1)\frac{n}{\sqrt{|P|}}) = n\sqrt{|P|} + |P|\frac{n}{\sqrt{|P|}} + n = n(2\sqrt{|P|}+1) \ . \quad \square$$

We now use Lemmas 6 and 7 to prove the upper bound.

Theorem 5 For all $k \geq 1$,

$$\nu_{1,k} \leq \frac{c_2}{k^2} \ .$$

Proof: Let $P \subseteq V_n$ be $(1,k)$-vertex-disjoint embedding and let W be a chorded k-cycle (or a chorded $(k-1)$-cycle if k is odd) with crossing number $\geq k^2/10^8$. Such a chorded cycle exists by Lemma 6. Denote the graph whose edges are the chords of W by W_c. Label the vertices of the chorded cycle $v_1, \ldots, v_k$, sequentially on the cycle. Let $p_1, \ldots, p_k$ be a sequence of elements of P such that there is a simple path of length at most $\frac{4.2kn}{\sqrt{|P|}}$ connecting $p_1, \ldots, p_k$ in that order. Lemma 7 guarantees the existence of such a sequence. Complete the path to a cycle in the plane by adding a parallel path alongside it and connecting the corresponding ends. The length of this cycle is at most $\frac{8.4kn}{\sqrt{|P|}}$. Identify vertex v_i of W with vertex p_i of V_n and find a vertex-disjoint embedding of W_c in $V_n - P$. Such an embedding exists as P is $(1,k)$-vertex-disjoint embedding and W_c is a degree-1 graph of size k.

Next, connect the simple path that goes through $p_1, \ldots, p_k$. The resulting graph is isomorphic to W and hence has a crossing number of at least $k^2/10^8$. These crossings are not cycle to cycle or chord to chord. Therefore they are all cycle edges crossing chord edges. The distance between any two chords is at least a grid unit. Therefore the total length of the path is at least $k^2/10^8$. Thus, $k^2/10^8 \leq$ length of the path $\leq \frac{8.4kn}{\sqrt{|P|}}$, or

$$\frac{|P|}{n^2} \leq \frac{8 \cdot 10^{17}}{k^2} \ . \quad \square$$

A graph $G = (V,E)$ is *α-expanding* for $\alpha > 0$ if every subset $S \subseteq V$ of size at most $|V|/2$ is adjacent to at least $\alpha|S|$ vertices in $V - S$. We now prove Lemma 6. The proof uses the following interesting result proved in [Le 84]. (The constant 97 is derived using the planar-separator bound of [Ch 88].)

Lemma 8 [Le 84] The crossing number of a k-node, α-expanding graph is at least $k^2\left[\left(\frac{\alpha}{10+\alpha}\right)^2 \frac{1}{97} - \frac{1}{k}\right]$. $\square$

We show that there exists a $\frac{1}{100}$-expanding chorded k-cycle. Let C_k be the k-vertex circle-graph with vertices: $\{1, \ldots, k\}$ and edges: $\{\{k,1\}, \{1,2\}, \{2,3\}, \ldots, \{k-1,k\}\}$. A nonempty set $I \subseteq \{1, \ldots, k\}$ is an *island of* C_k if the graph $C_k|_I$ is connected. Every set $S \subseteq \{1, \ldots, k\}$ decomposes uniquely into islands of C_k. These are the connected components of $C_k|_S$. For example, the set $\{1,3\}$ consists of 2 islands, $\{1\}$ and $\{3\}$, in C_4 and of the single island $\{1,3\}$ in C_3.

Lemma 9 The number of sets of size s that consist of i islands in C_k is

$$\binom{s}{i}\binom{k-s-1}{i-1}+\binom{s-1}{i-1}\binom{k-s}{i}=\frac{k}{i}\binom{s-1}{i-1}\binom{k-s-1}{i-1}\ .$$

Proof: We use the following facts:

1. The number of ways to partition the circle $1,\ldots,l$ into m islands is $\binom{l}{m}$.

2. The number of ways to partition the circle $1,\ldots,l$ into m islands so that 1 and l belong to different islands is $\binom{l-1}{m-1}$.

Each set S of size s consisting of i islands in C_k that contains the point 1 uniquely corresponds to a partition of the circle $1,\ldots,s$ into i islands combined with a partition of the circle $1,\ldots,k-s$ into i islands so that 1 and $k-s$ are in different islands. Hence there are $\binom{s}{i}\binom{k-s-1}{i-1}$ such sets.

Similarly, each size s set consisting of i islands in C_k that does not contain the point 1, uniquely corresponds to a partition of the circle $1,\ldots,s$ into i islands so that 1 and s are in different islands and a partition of the circle $1,\ldots,k-s$ into i islands. Therefore, there are $\binom{s-1}{i-1}\binom{k-s}{i}$ such sets. □

Example 2 Lemma 9 implies that there are $\binom{5}{3}\binom{3}{2}+\binom{4}{2}\binom{4}{3}=54$ sets of size 5 consisting of 3 islands in C_9. They are all 9 cyclic shifts of the sets in Figure 16. □

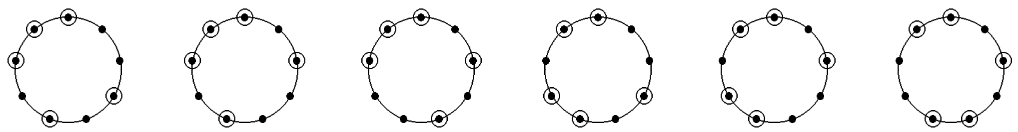

Figure 16: Sets of size 5 consisting 3 islands in C_9.

Let k be even, $S\subseteq\{1,\ldots,k\}$ a fixed set of size s, and $t<s$. Consider all $\prod_{i=1}^{k/2}(k+1-2i)$ different chorded k-cycles. If s and t have different parities, no chorded k-cycle has exactly t chords connecting elements of S to elements of $\{1,\ldots,k\}-S$. If s and t have the same parity, the number of chorded k-cycles with exactly t spokes connecting S to $\{1,\ldots,k\}-S$ is

$$\binom{s}{t}\cdot\prod_{i=1}^{t}(k-s+1-i)\cdot\prod_{i=1}^{\frac{s-t}{2}}(s-t+1-2i)\cdot\prod_{i=1}^{\frac{k-s-t}{2}}(k-s-t+1-2i)\ .$$

Therefore, if a chorded k-cycle is picked uniformly at random, the probability that t spokes connect S to $V-S$ is:

$$\binom{s}{t} \cdot \prod_{i=1}^{\lfloor t/2 \rfloor} \frac{k-s+2-(t)_2-2i}{k+1-2i} \cdot \prod_{i=1}^{\frac{s-t}{2}} \frac{s-t+1-2i}{k-t+1+(t)_2-2i}$$

where $(t)_2$ denotes the remainder of t when divided by 2. Thus the probability that S is connected to $\{1,\ldots,k\}-S$ by fewer than αs spokes, $\alpha < \frac{1}{2}$, is bounded above by

$$\sum_{t=0}^{\lceil \alpha s \rceil -1} \binom{s}{t} \cdot \prod_{i=1}^{\frac{s-t}{2}} \frac{s-t+1-2i}{k-t+1-2i} \leq \lceil \alpha s \rceil \binom{s}{\lceil \alpha s \rceil -1} \cdot \prod_{i=1}^{\frac{s-\lceil \alpha s \rceil+1}{2}} \frac{s-\lceil \alpha s \rceil+2-2i}{k-\lceil \alpha s \rceil+2-2i} .$$

If $0 < s \leq \frac{k}{2}$ and S consists of $i < s$ islands, then at least $i+1$ elements of $\{1,\ldots,k\}-S$ are connected to elements of S via cycle edges. Therefore, any set of size $0 < s \leq \frac{k}{2}$ connected to fewer than αs elements of $\{1,\ldots,k\}-S$ consists of at most $\lceil \alpha s \rceil - 2$ islands. From the last lemma, there are at most $\sum_{i=1}^{\lceil \alpha s \rceil -2} \left(\binom{s}{i}\binom{k-s-1}{i-1} + \binom{s-1}{i-1}\binom{k-s}{i}\right)$ such sets. For $\alpha \leq \frac{1}{2}$, this is at most $(\lceil \alpha s \rceil -2)\binom{s}{\lceil \alpha s \rceil -1}\binom{k-s}{\lceil \alpha s \rceil -1}$.

Hence, for $\alpha < \frac{1}{3}$, the probability that a randomly chosen set of size $s \leq \frac{k}{2}$ will be connected to its complement by fewer than αs edges is at most:

$$\begin{aligned}
&\lceil \alpha s \rceil (\lceil \alpha s \rceil -2)\binom{s}{\lceil \alpha s \rceil -1}^2 \binom{k-s}{\lceil \alpha s \rceil -1} \prod_{i=1}^{\frac{s-\lceil \alpha s \rceil+1}{2}} \frac{s-\lceil \alpha s \rceil+2-2i}{k-\lceil \alpha s \rceil+2-2i} \\
&\leq \alpha^2 s^2 \binom{s}{\lceil \alpha s \rceil -1}^2 2^{\lceil \alpha s \rceil -1} \prod_{i=1}^{\frac{s-3\lceil \alpha s \rceil+3}{2}} \frac{s-\lceil \alpha s \rceil+2-2i}{k-\lceil \alpha s \rceil+2-2i} \\
&\leq \alpha^2 s^2 \binom{s}{\lceil \alpha s \rceil -1}^2 2^{\alpha s} \left(\frac{s-\lceil \alpha s \rceil}{k-\lceil \alpha s \rceil}\right)^{\frac{s-3\lceil \alpha s \rceil+3}{2}} \\
&\leq \frac{\alpha s}{2\pi(1-\alpha)} 2^{s\left(2h(\alpha)+\frac{5}{2}\alpha-\frac{1}{2}\right)} ,
\end{aligned}$$

where we used the inequalities

$$\binom{k-s}{\lceil \alpha s \rceil -1} \prod_{i=\frac{s-3\lceil \alpha s \rceil+5}{2}}^{\frac{s-\lceil \alpha s \rceil+1}{2}} \frac{s-\lceil \alpha s \rceil+2-2i}{k-\lceil \alpha s \rceil+2-2i} \leq 2^{\lceil \alpha s \rceil -1} ,$$

and (cf.[Ga 68])

$$\binom{s}{t} < \sqrt{\frac{s}{2\pi t(s-t)}} \cdot 2^{s \cdot h(\frac{t}{s})} .$$

For $\alpha = .01$ we obtain that the probability there is a set of size s that is connected to fewer than $\frac{s}{100}$ elements in its complement is at most $\frac{s}{100}2^{-s/3}$. Therefore, the probability that a chorded-cycle is not an expander is bounded above by $\frac{1}{100}\sum_{s=1}^{k/2} s2^{-s/3} < 1$. We have proved:

Lemma 10 For all even k, there is a $\frac{1}{100}$-expander chorded k-cycle. $\square$

Lemmas 8 and 10 prove Lemma 6 and hence Theorem 5.

7 Alternative Models

We now discuss some variations of the PGA model and their effect on the achievable density. We heretofore assumed that the electrical elements are placed in the processing nodes. In most real PGA's they are placed in the 'squares' between the wires. This is modeled by the modified PGA graph shown in Figure 17(a).

Arguments similar to those used in Sections 5 and 6 can show that the same asymptotic results hold for such a PGA, too. Even if each processing element were connected to all of its four neighboring nodes, as shown in Figure 17(b), the asymptotic densities would not change.

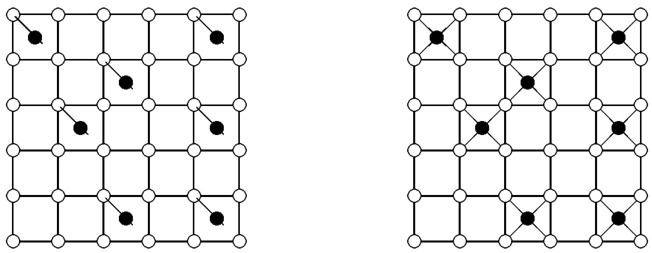

(a) (b)

Figure 17: PGA's with processors in the squares.

We assumed that a processing node cannot connect two wires even if the node is not used in the specific implementation. This is the most restrictive of several possible assumptions. The least stringent assumption allows a processing node to connect two wires even when the element it contains is used.

Relaxing this assumption increases the highest achievable density in both vertex and edge connectivity. For example, $\nu_{1,2}$ increases from 2/3 to 1. Yet the asymptotic results remain valid. This is clear in edge connectivity. To see that it holds in vertex connectivity, consider the chorded k-cycle with crossing number $\Omega(k^2)$ in the proof of Theorem 5. At most k of the crossings can correspond to edges passing through vertices of the chorded k-cycle. Still, there must be $\Omega(k^2)$ crossings between chords and rim edges. Hence the proof can proceed as before.

Allowing processing nodes to connect wires will not make a difference in the graph of Figure 17(a) since each processing node is connected to just one edge. Allowing the processing nodes of Figure 17(b) to connect wires will result in a PGA graph similar to that of Figure 18. This graph is discussed next.

Consider the graph of Figure 18. We show that any grid PGA with edge-connecting nodes can be simulated by a PGA based on the graph of Figure 18 with vertex-connecting nodes with only a factor of four loss in density. Associate node (i, j) of the grid PGA with node $(2i, 2j)$ of the new PGA and leave all other nodes as switching nodes. It is not hard to see that each switching node of the original PGA together with three of its new neighbors can simulate an edge-connecting switching node. This is interesting because the PGA of Figure 18 uses only twice the number of wires used in a standard grid PGA. Yet even with *vertex* connectivity, it allows for a $\Theta\left(\frac{1}{k}\right)$ component density; viz., the same density which is achievable only with *edge* connectivity in standard grid PGAs. This occurs partly because the current PGA graph is not planar.

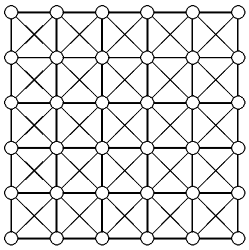

Figure 18: PGA consisting of the grid and the diagonal edges.

We assumed that each edge of the grid represents one wire. In practice, each edge corresponds to a *channel* consisting of several wires. This enables implementation of circuits with degree larger than four. Allowing a fixed number of wires in every channel increases the achievable density by a constant factor but the asymptotic results remain valid.

Perhaps the most critical departure of our model from real PGA's is that in practice each component represents a type of electrical element such as a gate or a memory cell. A real PGA contains many components of each type. A circuit design specifies the type of component that must be used and every component of that type will do.

One extreme version of the problem assumes that all components are of the same type. This is the traditional graph embedding problem discussed in Chapter 3 of [Ul 83] and the references therein. The user has the freedom to choose the mapping both of the vertices and of the edges. We assumed the other extreme, namely that every component is of a different type. Hence the vertex mapping is given. The intermediate problem where there is a certain number of components of each type is interesting but was not dealt with here. So is the problem of realizing all circuits of a given size with a large PGA having only one type of component.

Acknowledgements

We are indebted to the referee who improved our upper bound on $\nu_{1,4}$ from 6/11 to 1/2, thereby establishing its exact value. We also thank Peter Shor for acquainting us with [AKLLW 85] and [Ka 89].

References

[AKLLW 85] A. Aggarwal, M. Klawe, D. Lichtenstein, N. Linial, and A. Wigderson, 'Multi-Layer Grid Embeddings', Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1985, pp. 186-195.

[AKS 89] A. Aggarwal, M. Klawe, P. Shor, 'Multi-Layer Grid Embeddings for VLSI', to appear, Algorithmica.

[Ch 88] F.R.K. Chung, 'Improved Separators from Planar Graphs', preprint.

[CS 78] M. Cutler and Y. Shiloach, 'Permutation Layout', Networks, Vol. 8, 1978, pp. 253-278.

[EGRREM 88] A. El Gamal, J. Greene, J. Renyeri, E. Rogoyski, K. El-Ayat, A. Moshen, 'An Architecture for Electrically Configurable Gate Arrays', IEEE 1988 Custom Integrated Circuits Conference, pp. 15.4.1-15.4.5.

[Fe 55] L. Few, 'The Shortest Path and Shortest Road Through n Points', Mathematika 2, 1955, pp. 141-144.

[Fr 88] R. Freeman, 'User-Programmable Gate Arrays', IEEE Spectrum, December 1988, pp. 32-35.

[Ga 68] R.G. Gallager, *Information Theory and Reliable Communication*, John Wiley and Sons, 1968.

[Ka 89] H. Karloff, 'How Long Can a Euclidean Traveling Salesman Tour Be', Siam Journal of Discrete Math., Vol. 2, No. 1, pp. 91-99.

[Le 84] F.T. Leighton, 'New Lower Bound Techniques for VLSI', Math Systems Theory 17, 1984, pp. 47-70.

[Ul 83] J.D. Ullman, *Computational Aspects of VLSI*, Computer Science Press, 1983.