Geometric Thickness in a Grid of Linear Area*

David R. Wood

Basser Department of Computer Science
The University of Sydney
Sydney NSW 2006, Australia
davidw@cs.usyd.edu.au

Let $G = (V, E)$ be a connected graph with $n = |V|$ vertices, $m = |E|$ edges, maximum degree Δ, and genus γ. The *thickness* of G, denoted by $\theta(G)$, is the minimum number of subgraphs in a partition of the edges into planar subgraphs. See [MOS98] for a survey of results concerning this well-studied parameter. The *geometric thickness* of G, denoted by $\overline{\theta}(G)$, is the minimum number of layers such that G can be drawn in the plane with edges as straight-line segments, and with each edge assigned to a layer so that no two edges in the same layer cross. Geometric thickness was first introduced under the name of *real linear thickness* [Kai73]. Lower and upper bounds are established for the geometric thickness of complete and complete bipartite graphs in [DEH00]. In particular, it is shown that $\overline{\theta}(K_n) > \theta(K_n)$ for large n. On the other hand, $\overline{\theta}(K_{a,b}) = \theta(K_{a,b})$ when $a \gg b$. Applications of geometric thickness include multilayer VLSI [AKS91] and graph visualisation [DETT99].

Another parameter closely related to geometric thickness is that of book thickness. A *book* consists of a line in 3-space, called the *spine*, and some number of *pages*, each a half-plane with the spine as boundary. A *book embedding* of a graph consists of a linear ordering of the vertices along the spine of a book and an assignment of edges to pages so that edges assigned to the same page can be drawn on that page without crossings. The *book thickness* or *pagenumber* of a graph G, denoted by $\mathsf{bt}(G)$, is the minimum number of pages in a book embedding of G. Book thickness is equivalent to the thickness with the additional requirement that the subgraphs are outerplanar. To see this, position the vertices on a circle in the plane according to their order along the spine. Furthermore, the edges can be straight line-segments in such a drawing [BK79]. Thus, the book thickness of a graph is the same as its 'geometric outerplanar thickness', and thus, $\overline{\theta}(G) \leq \mathsf{bt}(G)$.

Other upper bounds for the thickness of arbitrary graphs are $\lfloor \sqrt{m/3} + 3/2 \rfloor$ [DHS91], $6 + \sqrt{2\gamma - 2}$ [DH91], and $\lceil \Delta/2 \rceil$ [Hal91]. Of these three upper bounds, the first two are asymptotically matched by the book thickness. That is, $\mathsf{bt}(G) \in O(\sqrt{\gamma})$ [Mal94a], and since $\gamma \leq m$, $\mathsf{bt}(G) \in O(\sqrt{m})$, a result proved independently in [Mal94b]. It is an open problem to decide if the third of these upper bounds is asymptotically matched by the geometric thickness; that is, is $\overline{\theta}(G) \in O(\Delta)$?

A graph G has a drawing in the plane with an arbitrary set of preassigned vertex locations with edges assigned to $\theta(G)$ layers such that no two edges in the same layer cross [Kai73]. Thus, the key difference between geometric thickness and graph-theoretic thickness is that geometric thickness requires the edges to be drawn as straight line-segments, whereas the graph-theoretic thickness allows edges to bend arbitrarily.

In this paper we consider a variation of geometric thickness which lies between thickness and geometric thickness in which each edge has at most one bend. We are also interested in drawings with small area, which is an important consideration in VLSI and visualisation. To measure the area of a drawing we assume a *vertex resolution rule*; that is, pairs of vertices are at least unit-distance apart. A drawing obtained from a book embedding by positioning the vertices around a circle, as discussed above, has $O(n^2)$ area. The construction in [DEH00] demonstrating that $\overline{\theta}(K_n) \leq \lceil \frac{n}{4} \rceil$ has $O(n^6)$ area [D. Eppstein, personal communication]. We prove the following 2-dimensional generalisation of the above-mentioned result in [Mal94b] for producing book embeddings.

Theorem 1. *The vertices of G can be positioned in a $2\sqrt{n} \times 2\sqrt{n}$ grid, with each edge drawn with at most one bend, such that the edges can be partitioned into $O(\sqrt{m})$ layers, and no two edges in the same layer cross. There is a Las Vegas algorithm to produce the drawing in $O(m \log^3 n \log\log n)$ time (with high probability).*

Proof Sketch. The $N \times N$ *grid* is the set of *grid points* $\{(x, y) \in \mathbb{N} \times \mathbb{N} : 1 \leq x, y \leq N\}$

* Supported by the Australian Research Council Large Grant A49906214. Completed while visiting the School of Computer Science, McGill University, Montréal.

in the plane. The x_0-*column* and the y_0-*row* of the $N \times N$ grid are the set of grid-points $\{(x_0, y) : 1 \le y \le N\}$ and $\{(x, y_0) : 1 \le x \le N\}$, respectively. Assuming N is a power of 2, for $a = 1, 2, \ldots, \log N$, partition the grid-points into 2^a sets of consecutive columns of equal size, called the *vertical a-strips*, which we label L, R, L, R, etc. from left to right. For $b = 1, 2, \ldots, \log N$, partition the grid-points into 2^b sets of consecutive rows of equal size, called the *horizontal b-strips*, which we label D, U, D, U, etc. from bottom to top. The intersection of a vertical a-strip and a horizontal b-strip is called an (a, b)-*region*. Note that a $(\log N, b)$-region is contained within a single column, and an $(a, \log N)$-region is contained within a single row. An (a, b)-region S labelled L and D is *coupled* with the (a, b)-region S' labelled R and U immediately above and to the right of S. An (a, b)-region S labelled R and D is *coupled* with the (a, b)-region S' labelled L and U immediately above and to the left of S. In addition, a $(\log N, b)$-region S labelled D is *coupled* with the $(\log N, b)$-region S' labelled U immediately above it. Similarly, an $(a, \log N)$-region S labelled L is *coupled* with the $(a, \log N)$-region S' labelled R immediately to the right it.

An $N \times N$ *grid layout* is a 1-1 mapping π from the vertices of a graph to the grid-points of an $N \times N$ grid. Denote $\pi(v)$ by $(X_\pi(v), Y_\pi(v))$ for each vertex v. Assuming G has $n = N^2$ vertices, let π be a random grid layout of G in the $N \times N$-grid. (The $N \times N$ grid layouts of G correspond to the permutations of the vertices.)

For each edge vw, there is a unique coupled pair of regions with one region containing v and the other region containing w (see Fig. 1). We say vw *belongs* to this couple. For a given coupled pair of regions, the (2-coloured bipartite) subgraph of edges belonging to it is called a *regional* subgraph.

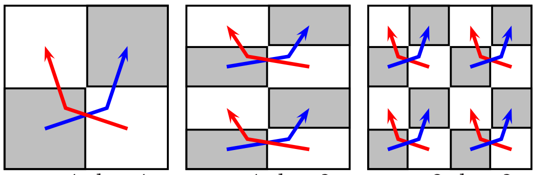

Fig. 1. (a, b)-regions of a grid layout.

Partition and orient the edges as follows. For each x_0, y_0, $1 \le x_0, y_0 \le N$, let $E^+_{x_0,y_0} = \{\overrightarrow{vw} \in E : X_\pi(v) < X_\pi(w) = x_0$ and $Y_\pi(v) = y_0 < Y_\pi(w)\}$, $E^-_{x_0,y_0} = \{\overrightarrow{vw} \in E : Y_\pi(v) = y_0 < Y_\pi(w)$ and $X_\pi(v) > X_\pi(w) = x_0\}$, $E^r_{y_0} = \{\overrightarrow{vw} \in E : Y_\pi(v) = Y_\pi(w) = y_0$ and $X_\pi(v) < X_\pi(w)\}$, $E^c_{x_0} = \{\overrightarrow{vw} \in E : X_\pi(v) = X_\pi(w) = x_0$ and $Y_\pi(v) < Y_\pi(w)\}$. Edges in these four classes respectively belong to the four types of regional subgraphs. We use different sets of layers for edges in each of these four classes, while edges in one class and belonging to coupled pairs of (a, b)-regions (for fixed a and b) share the same set of layers.

For each $E^\pm_{x_0,y_0}$, assign evenly-spaced bends along the line-segment from $\left(x_0 \mp \frac{1}{2},\, y_0 + \frac{1}{2}\right)$ to $\left(x_0 \mp \frac{1}{2} \mp \frac{1}{\sqrt{2}\,N},\, y_0 + \frac{1}{2} + \frac{1}{\sqrt{2}\,N}\right)$ to the edges $\overrightarrow{vw} \in E^\pm_{x_0,y_0}$ in increasing order of $|x_0 - X_\pi(v)| + |y_0 - Y_\pi(w)|$ (breaking ties arbitrarily); see Fig. 2.

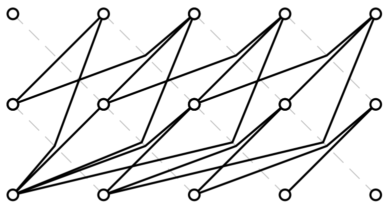

Fig. 2. Edges in E^+.

An edge $\overrightarrow{vw} \in E^r_{y_0}$ is drawn with a bend at $\left(\frac{1}{2}\left(X_\pi(w) - X_\pi(v)\right),\, Y_\pi(v) + \left(\frac{1}{N}(X_\pi(w) - X_\pi(v) - 1)\right)^2\right)$, and an edge $\overrightarrow{vw} \in E^c_{x_0}$ is drawn with a bend at $\left(\frac{1}{2}\left(Y_\pi(w) - Y_\pi(v)\right),\, X_\pi(v) + \left(\frac{1}{N}(Y_\pi(w) - Y_\pi(v) - 1)\right)^2\right)$; see Fig. 3. (This method essentially describes how to draw a book embedding in $O(n)$ area with at most one bend per edge.)

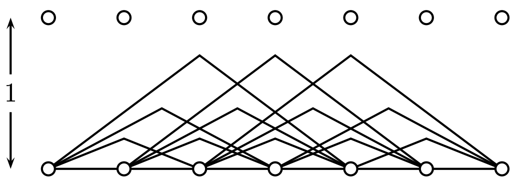

Fig. 3. Edges in E^r.

Edge crossings can be characterised. A set of pairwise crossing edges contained in a regional subgraph are said to be *completely crossing*. Since edges incident to a common vertex do not cross, a completely crossing set of edges is a matching. The following lemma generalises

Lemma 2.2 in [Mal94b], and is proved by applying Dilworth's Theorem [Dil50] to a certain partial order.

Lemma 1. *The edges of a 2-coloured regional subgraph H of G with at most k completely crossing edges can be partitioned into k layers in $O(m \log\log n)$ time, such that with the edges drawn with one bend as described above, no two edges in a single layer cross.* □

We group regions by the number of grid-points they contain. An (a,b)-region has $\frac{n}{2^{a+b}}$ grid-points. The regional subgraph associated with a coupled pair of (a,b)-regions with $a+b=j$ is said to be in the *j-level*. For $2 \le j \le 2\log N$, let A_k^j be the event that there exists a 2-coloured matching M in G with k edges such that M is contained in a j-level regional subgraph and is completely crossing. The probability of A_k^j is

$$\mathbf{P}\left\{A_k^j\right\} < \underbrace{\binom{m}{k}2^k}_{(1)} \cdot \underbrace{(j+2)2^{j-1}}_{(2)} \cdot \underbrace{\binom{\frac{n}{2^j}}{k}^2 \frac{(k!)^2(n-2k)!}{n!}}_{(3)} \cdot \underbrace{\frac{1}{k!}}_{(4)},$$

where:

(1) is an upper bound on the number of 2-coloured matchings M with k edges,
(2) is an upper bound on the number of j-level regional subgraphs,
(3) is the probability that every edge in M belongs to a fixed j-level subgraph, and
(4) is the probability that M is completely crossing in π, given fixed grid-points for the end-vertices of M.

Applying well known formulae and other observations, $\mathbf{P}\{A_k^j\} < (\frac{1}{k}\mathbf{e}^2(j+2)\sqrt{m/2^{j+1}})^{2(k+1)}$. Define $k_j = \mathbf{e}^2(j+2)\sqrt{\frac{m}{2^j}}$. Thus, $\mathbf{P}\{A_{k_j}^j\} < 2^{-1-\mathbf{e}^2(j+2)\sqrt{m/2^j})} < 2^{-1-2\mathbf{e}^2\sqrt{n/2^j}}$. The probability that $A_{k_j}^j$ does not occur for all j, $2 \le j \le 2\log N$, is at least $1-\sum_j 2^{-1-2\mathbf{e}^2\sqrt{n/2^j}} > 1-2^{-2\mathbf{e}^2} > 0.99996$.

Thus, there exists a grid layout π' of G such that $A_{k_j}^j$ does not occur for all j. That is, in each regional j-level subgraph, there is no completely crossing matching with at least k_j edges (with respect to π'). Applying Lemma 1 for each regional subgraph in π', the j-level edges take up at most $2(j+2)k_j$ layers, and the total number of layers is at most $\sum_j 2(j+2)k_j = 2\mathbf{e}^2\sqrt{m}\sum_j(j+2)^2 2^{-j/2} \le (2+\sqrt{2})(13\sqrt{2}+31)\mathbf{e}^2\sqrt{m} < 1246\sqrt{m}$.

By iterating the above algorithm at most $\log n$ times (choosing a random grid layout in each iteration), the probability of failure at every iteration is at most $n^{-2\mathbf{e}^2}$. The time bound for our Las Vegas algorithm follows. This completes the sketch of the proof. □

References

[AKS91] A. Aggarwal, M. Klawe, and P. Shor. Multilayer grid embeddings for VLSI. *Algorithmica*, 6(1):129–151, 1991.

[BK79] F. Bernhart and P. C. Kainen. The book thickness of a graph. *J. Combin. Theory Ser. B*, 27(3):320–331, 1979.

[DEH00] M. B. Dillencourt, D. Eppstein, and D. S. Hirschberg. Geometric thickness of complete graphs. *J. Graph Algorithms Appl.*, 4(3):5–17, 2000.

[DETT99] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs.* Prentice-Hall, New Jersey, 1999.

[DH91] A. M. Dean and J. P. Hutchinson. Relations among embedding parameters for graphs. In *Graph theory, combinatorics, and applications, Vol. 1 (Kalamazoo, MI, 1988)*, pages 287–296. Wiley, New York, 1991.

[DHS91] A. M. Dean, J. P. Hutchinson, and E. R. Scheinerman. On the thickness and arboricity of a graph. *J. Combin. Theory Ser. B*, 52(1):147–151, 1991.

[Dil50] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Ann. of Math. (2)*, 51:161–166, 1950.

[Hal91] J. H. Halton. On the thickness of graphs of given degree. *Inform. Sci.*, 54(3):219–238, 1991.

[Kai73] P. C. Kainen. Thickness and coarseness of graphs. *Abh. Math. Sem. Univ. Hamburg*, 39:88–95, 1973.

[Mal94a] S. M. Malitz. Genus g graphs have pagenumber $O(\sqrt{g})$. *J. Algorithms*, 17(1):85–109, 1994.

[Mal94b] S. M. Malitz. Graphs with E edges have pagenumber $O(\sqrt{E})$. *J. Algorithms*, 17(1):71–84, 1994.

[MOS98] P. Mutzel, T. Odenthal, and M. Scharbrodt. The thickness of graphs: a survey. *Graphs Combin.*, 14(1):59–73, 1998.