

Densities of Minor-Closed Graph Families

David Eppstein

Computer Science Department
University of California, Irvine
Irvine, California, USA

Submitted: Jul 1, 2010; Accepted: Sep 22, 2010; Published: Oct 15, 2010
Mathematics Subject Classification: 05C83 (Primary) 05C35, 05C42 (Secondary)

Abstract

We define the limiting density of a minor-closed family of simple graphs \mathcal{F} to be the smallest number k such that every n -vertex graph in \mathcal{F} has at most $kn(1 + o(1))$ edges, and we investigate the set of numbers that can be limiting densities. This set of numbers is countable, well-ordered, and closed; its order type is at least ω^ω . It is the closure of the set of densities of density-minimal graphs, graphs for which no minor has a greater ratio of edges to vertices. By analyzing density-minimal graphs of low densities, we find all limiting densities up to the first two cluster points of the set of limiting densities, 1 and $3/2$. For multigraphs, the only possible limiting densities are the integers and the superparticular ratios $i/(i + 1)$.

1 Introduction

Planar simple graphs with n vertices have at most $3n - 6$ edges. Outerplanar graphs have at most $2n - 4$ edges. Friendship graphs have $3(n - 1)/2$ edges. Forests have at most $n - 1$ edges. Matchings have at most $n/2$ edges. Where do the coefficients 3, 2, $3/2$, 1, and $1/2$ of the leading terms in these bounds come from?

Planar graphs, forests, outerplanar graphs, and matchings all form instances of *minor-closed families of simple graphs*, families of the graphs with the property that any *minor* of a graph G in the family (a simple graph formed from G by contracting edges and removing edges and vertices) remains in the family. The friendship graphs (graphs in the form of $(n - 1)/2$ triangles sharing a common vertex) are not minor-closed, but are the maximal graphs in another minor-closed family, the graphs formed by adding a single vertex to a matching. For any minor-closed family \mathcal{F} of simple graphs there exists a number k such that every n -vertex graph in \mathcal{F} has at most $kn(1 + o(1))$ edges [7, 18, 19]. We define the *limiting density* of \mathcal{F} to be the smallest number k with this property. In other words, it is the coefficient of the leading linear term in the *extremal function* of \mathcal{F} , the function that maps a number n to the maximum number of edges in an n -vertex

graph in \mathcal{F} . We may rephrase our question more formally, then, as: which numbers can be limiting densities of minor-closed families?

As we show, the set of limiting densities is countable, well-ordered, and topologically closed (Theorem 19). Additionally, the set of limiting densities of minor-closed graph families is the closure of the set of densities of a certain family of finite graphs, the *density-minimal graphs* for which no minor has a greater ratio of edges to vertices (Theorem 20). To prove this we use a separator theorem for minor-closed families [1] to find density-minimal graphs that belong to a given minor-closed family, have a repetitive structure that allows them to be made arbitrarily large, and are close to maximally dense.

By analyzing the structure of density-minimal graphs, we can identify the smallest two cluster points of the set of limiting densities, the numbers 1 and $3/2$, and all of the other possible limiting densities that are at most $3/2$. Specifically, the limiting densities below 1 are the *superparticular ratios* $i/(i+1)$ for $i = 0, 1, 2, \dots$, and the corresponding density-minimal graphs are the $(i+1)$ -vertex trees. The limiting densities between 1 and $3/2$ are the rational numbers of the form $3i/(2i+1)$, $(3i+2)/(2i+2)$, and $(3i+4)/(2i+3)$ for $i = 1, 2, 3, \dots$; the corresponding density-minimal graphs include the friendship graphs and small modifications of these graphs. Beyond $3/2$ the pattern is less clear, but each number $2 - 1/i$ is a cluster point in the set of limiting densities, and 2 is a cluster point of cluster points. Similarly the number 3 is a cluster point of cluster points of cluster points, etc. We summarize this structure in Theorem 22.

One may also apply the theory of graph minors to families of *multigraphs* allowing multiple edges between the same pair of vertices as well as multiple self-loops connecting a single vertex to itself. In this case the theory of limiting densities and density-minimal graphs is simpler: the only possible limiting densities for minor-closed families of multigraphs are the integers and the superparticular ratios, and the only limit point of the set of limiting densities is the number 1 (Theorem 23).

2 Related work

2.1 Limiting densities of minor-closed graph families

A number of researchers have investigated the limiting densities of minor-closed graphs. It is known that every minor-closed family has bounded limiting density [11] and that this density is at most $O(h\sqrt{\log h})$ for graphs with an h -vertex forbidden minor [7, 18]. This asymptotic growth rate is tight: K_h -free graphs have limiting density $\Theta(h\sqrt{\log h})$, and the constant factor hidden in the Θ -notation above is known [19].

There have also been similar investigations into the dependence of the limiting density of H -minor-free graphs on the number of vertices in H when H is not a clique [8, 12]. However these works have a different focus than ours: they concern either the asymptotic growth rate of the limiting density as a function of the forbidden minor size or, in some cases, the limiting densities or extremal functions of specific minor-closed families [4, 9, 15, 17] rather than, as here, the structure of the set of possible limiting densities.

2.2 Growth rates of minor-closed graph families

Bernardi, Noy, and Welsh [3] investigate a different set of real numbers defined from minor-closed graph families, their growth rates. The growth rate of a family of graphs is a number c such that the number of n -vertex graphs in the family, with vertices labeled by a permutation of the numbers from 1 to n , grows asymptotically as $n!c^{n+o(n)}$ [14]. Bernardi et al. investigate the topological structure of the set of growth rates, show that this set is closed under the doubling operation, and determine all growth rates that are at most 2.25159.

2.3 Upper density of infinite graphs

For arbitrary graphs, not belonging to a minor-closed family, the density is often defined differently, as the ratio

$$|E|/\binom{|V|}{2} = \frac{2|E|}{|V|(|V| - 1)}$$

of the number of edges that are present in the graph to the number of positions where an edge could exist. (This definition is not useful for minor-closed families: for any nontrivial minor-closed family, the density defined in this way necessarily approaches zero in the limit of large n .)

This definition of density can be extended to infinite graphs as the *upper density*: the upper density of an infinite graph G is the supremum of numbers α with the property that G contains arbitrarily large subgraphs with density α . Although defined in a very different way to our results here, the set of possible upper densities is again limited to a well-ordered countable set, consisting of 0, 1, and the superparticular ratios $i/(i + 1)$ [5, Exercise 12, p. 189].

3 Density-minimal graphs

We define the *density* of a simple graph G , with m edges and n vertices, to be the ratio m/n . We say that G is *density-minimal* if no proper minor of G has equal or greater density. Equivalently, a connected graph G is density-minimal if there is no way of contracting some of the edges of G , compressing multiple adjacencies to a single edge, and removing self-loops, that produces a smaller graph with greater density: edge removals other than the ones necessary to form a simple graph are not helpful in producing dense minors. The *rank* of a connected graph G , again with m edges and n vertices, is the number $m + 1 - n$ of independent cycles in G ; we say that G is *rank-minimal* if no proper minor of G has the same rank. Every density-minimal graph is also rank-minimal, because a smaller graph with equal rank would have greater density.

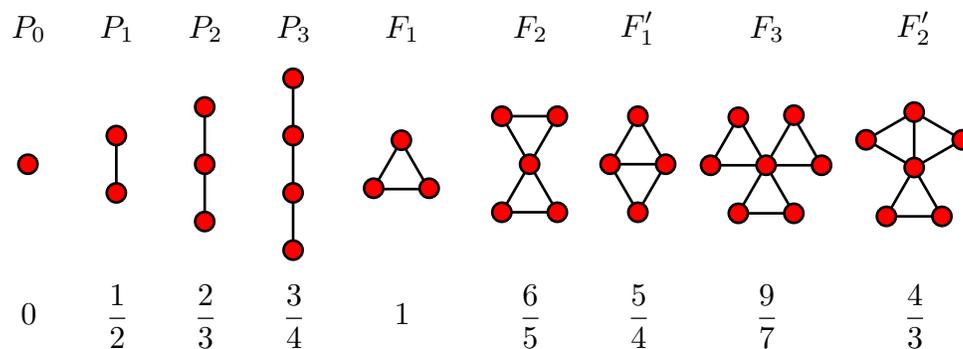


Figure 1: Some density-minimal graphs and their densities.

3.1 Examples of density-minimal graphs

Every tree is density-minimal: a tree with m edges (such as the path P_m) has density $m/(m+1)$, and contracting edges in a tree can only produce a smaller tree with a smaller density. Additional examples of density-minimal graphs are provided by the *friendship graphs* F_i formed from a set of i triangles by identifying one vertex from each triangle into a single supervertex; these graphs are famous as the finite graphs in which every two vertices have exactly one common neighbor [6], but they also have the property that F_i is density-minimal with density $3i/(2i+1)$. For, if we contract edges in F_i to form a minor H , then H must itself have the form of a friendship graph together possibly with some additional degree-one vertices connected to the central vertex. Removing the degree-one vertices can only increase the density of H , but once this is done H must itself be a friendship graph with fewer triangles than F_i and smaller density. This shows that every minor of F_i has smaller density, so F_i is density-minimal.

Let F'_i be formed from the friendship graph F_i by adding one more vertex, whose two neighbors are the two endpoints of any edge in F_i . Then a similar argument shows that F'_i is density-minimal with density $(3i+2)/(2i+2)$. If a second vertex is added in the same way to produce a graph F''_i , then F''_i is density-minimal with density $(3i+4)/(2i+3)$. However, adding a third vertex in the same way does not generally produce a density-minimal graph: its density is $3/2$, and (if $i > 2$) one of the triangles of the friendship graph from which it was formed can be removed leaving a smaller minor with the same density.

These are not the only density-minimal graphs with these densities—a set of $i \geq 3$ triangles can be connected together at shared vertices to form cactus trees other than the friendship graphs with the same density—but as we now show, their densities are the only possible densities of density-minimal graphs in this numerical range.

3.2 Bounding the rank of low-density density-minimal graphs

In order to determine the possible densities of density-minimal graphs, it is helpful to have the following technical lemma, which allows us to restrict our attention to graphs of

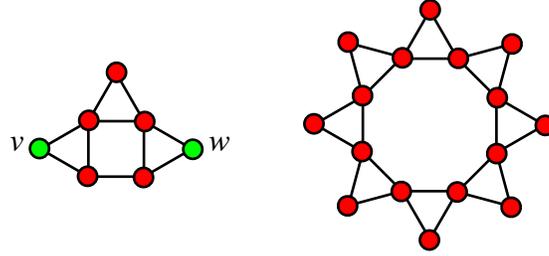


Figure 2: Left: a rank-four graph with seven vertices. After removing the two degree-two vertices v and w , the remaining graph G'' is a rank-two theta graph $\Theta(1, 2, 3)$; two of the edges on the length-three path of the theta are covered by the triangles containing v and w , but the remaining edge at the bottom of the drawing can be contracted to produce a smaller biconnected rank-four graph. Right: a cycle of triangles, an example of a high-rank biconnected graph in which all minors have density at most $3/2$.

low rank.

Lemma 1. *Every biconnected graph of rank four or higher contains a minor of density at least $3/2$.*

Proof. Let G be biconnected with rank four or higher. We perform an open ear decomposition of the graph (a partition of the edges into a sequence of subgraphs, the first of which is a cycle and the rest of which are simple paths, where the endpoints of each path belong to previous components of the decomposition) [20]. The first four ears of this decomposition form a biconnected subgraph of G with rank exactly four, so by replacing G with this subgraph we may assume without loss of generality that the rank of G is four. We may also assume without loss of generality that G has no edges that could be contracted in a way that preserves both its rank and its biconnectivity, because otherwise we could replace G with the graph formed by performing these contractions; in particular, this implies that every degree-two vertex of G is part of a triangle. Additionally, no two degree-two vertices can be adjacent, for if they were then the third vertex of their triangle would be an articulation point, contradicting biconnectivity.

We now assert that G has at most six vertices, and therefore has density at least $3/2$. For, suppose that G had n vertices with $n \geq 7$, and $n + 3$ edges. In this case, since the number of edges is less than $3n/2$, there must be a vertex v with degree two. Removing v leaves a smaller graph G' with $n - 1 \geq 6$ vertices and $n + 1$ edges; again, the number of edges is less than three-halves of the number of vertices, so there must be a vertex w that has degree two in G' (w cannot have degree one, because if it did then G would have two adjacent degree-two vertices). If w is not adjacent to v in G , then the two neighbors of w in G' form a triangle as they do in G , for the same reason that the neighbors of v form a triangle; if on the other hand w is adjacent to v , then its two remaining neighbors in G' must again form a triangle or else the edge wx where x is not adjacent to v could be contracted preserving rank and biconnectivity. Removing w from G' leaves a second smaller graph G'' with rank two. Additionally, since both v and w belonged to triangles

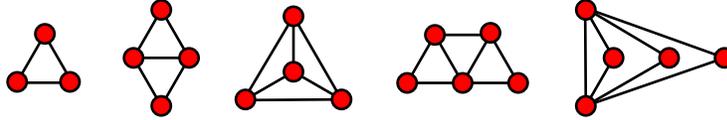


Figure 3: The five rank-minimal biconnected graphs with rank between one and three.

of G , their removal cannot create any articulation points in G'' , so G'' is biconnected.

But (by ear decomposition again) the only possible structure for a biconnected rank-two graph such as G'' is a theta graph $\Theta(a, b, c)$, in which two degree-three vertices are connected by three paths of lengths a , b , and c respectively. If G has seven or more vertices, then G'' has five or more vertices, and $a + b + c \geq 6$. If one of the paths of the theta graph has length three or more, then only two of the edges of this path can be part of triangles containing v and w , and the third edge of the path can be contracted in G to produce a smaller biconnected rank-four graph, contradicting our assumption that no such contraction exists; this case is shown in Figure 2(left). In the remaining case, $G'' = \Theta(2, 2, 2) = K_{2,3}$; only two of the six edges of G'' can be part of triangles involving v and w , and any one of the four remaining edges can be contracted in G to produce a smaller biconnected rank-four graph, again contradicting our assumption.

These contradictions show that the number n of vertices in G is at most six; since its rank is four, its density $(n + 3)/n$ must be at least $3/2$. \square

Increasing the rank past four without increasing the connectivity does not increase the density past the $3/2$ threshold of Lemma 1: there exist biconnected graphs of arbitrarily high rank in which the densest minors have density $3/2$, namely the cycles of triangles shown in Figure 2(right).

A simple case analysis based on ear decompositions shows that there are exactly five rank-minimal biconnected graphs of rank between one and three: the triangle $F_1 = K_3$, the diamond graph with four vertices and rank two, the complete graph K_4 with four vertices and rank three, and two different 2-trees with five vertices and rank three (Figure 3).

3.3 Classification of density-minimal graphs with low density

Lemma 2. *Let a graph G be density-minimal, with density $\Delta < 3/2$. Then*

$$\Delta \in \left\{ \frac{i}{i+1} \mid i \geq 0 \right\} \cup \left\{ \frac{3i+2j}{2i+j+1} \mid i \geq 1, j \in \{0, 1, 2\} \right\}$$

and for each number Δ in this set there exists a density-minimal graph G with density Δ .

Proof. First, suppose that $\Delta < 1$. Then G must be acyclic and connected, for if it had a cycle it would have a triangle minor with density 1 and not be density-minimal, and if it were disconnected then the densest of its components would be a minor with density at least as great as that of G itself. But an acyclic connected graph is a tree, and has density $i/(i + 1)$ where i is its number of edges.

In the remaining cases, $1 \leq \Delta < 3/2$. G must be bridgeless, because any bridge could be contracted producing a denser graph. By Lemma 1, each block (biconnected component) of G must have rank at most three, for otherwise that block by itself would have a minor with density at least $3/2$, contradicting the assumption that G is density-minimal. Additionally, each block must be rank-minimal, for otherwise G itself would not be rank-minimal. Therefore, each block must be a triangle, the diamond graph, or one of the two rank-three 2-trees. If there are two rank-three blocks, three rank-two blocks, or one rank-three block and one rank-two block, then those blocks alone (with the remaining blocks contracted away) would again form a minor with density at least $3/2$. The only remaining cases are a graph in which the blocks consist of i triangles, with density $3i/(2i + 1)$, a graph in which the blocks consist of i triangles and one rank-two block, with density $(3i + 5)/(2i + 4)$, a graph in which the blocks consist of i triangles and one rank-three block, with density $(3i + 7)/(2i + 5)$, or a graph in which the blocks consist of i triangles and two rank-two blocks, with density $(3i + 10)/(2i + 7)$. Each of these densities belongs to the set specified in the lemma.

To show that each Δ in the set of densities stated in the lemma is the density of some density-minimal graph G , we need only recall the path graphs P_i , the friendship graphs F_i , and the graphs F'_i and F''_i formed by adding one or two degree-two vertices to a friendship graph. As we have already argued, these graphs are density-minimal, and together they cover all the densities in the given set. \square

The set of achievable densities allowed by Lemma 2, in numerical order up to the limit point $3/2$, is

$$0, \frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{4}{5}, \frac{5}{6}, \dots, 1, \frac{6}{5}, \frac{5}{4}, \frac{9}{7}, \frac{4}{3}, \frac{15}{11}, \frac{11}{8}, \frac{18}{13}, \frac{7}{5}, \frac{24}{17}, \frac{17}{12}, \frac{27}{19}, \frac{10}{7}, \dots, \frac{3}{2}.$$

Figure 1 shows density-minimal graphs achieving some of these densities.

4 Fans of graphs

We now introduce a notation for constructing large graphs with a repetitive structure from a smaller model graph. Given a graph G , a proper subset S of the vertices of G , and a positive integer k , we define the graph $\text{Fan}(G, S, k)$ to be the union of k copies of G , all sharing the same copies of the vertices in S and having distinct copies of the vertices in $G \setminus S$. For instance:

- If G is a triangle uvw , then $\text{Fan}(G, \{u\}, k)$ is the friendship graph F_k
- For the same triangle $G = uvw$, $\text{Fan}(G, \{u, v\}, k)$ is a 2-tree formed by k triangles sharing a common edge. Three of the graphs in Figure 3 take this form, for $k \in \{1, 2, 3\}$.
- The complete bipartite graph $K_{a,b}$ can be represented in multiple different ways as a fan: it is isomorphic to $\text{Fan}(K_{d,b}, S, a/d)$ where d is any divisor of a and S is the

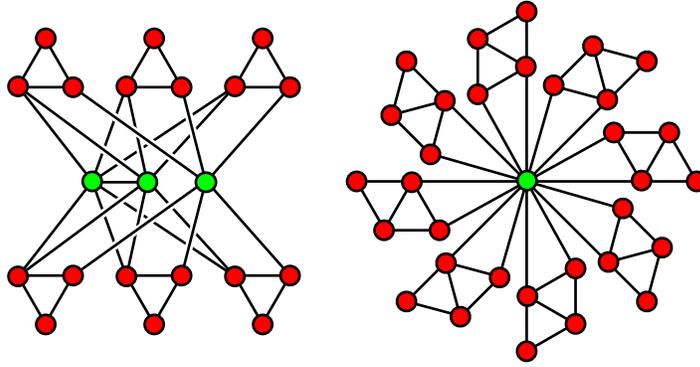


Figure 4: Two fans of graphs.

b -vertex side of the bipartition of $K_{d,b}$, and symmetrically there is a representation as a fan for any divisor of b .

- Two more examples are shown in Figure 4.

4.1 Basic observations about fans

If G has n vertices and S has s vertices (where $s < n$ as we require S to be a proper subset of the vertices), then $\text{Fan}(G, S, k)$ has $k(n - s) + s = \Omega(k)$ vertices. $\text{Fan}(G, \emptyset, k)$ is the disconnected graph formed by k disjoint copies of G ; however, if G is connected and S is nonempty then $\text{Fan}(G, S, k)$ is also connected.

Lemma 3. *Every graph $\text{Fan}(G, S, k)$ has at least k vertices.*

Proof. This follows immediately from the requirement that S be a proper subset of the vertices of G ; there is at least one vertex that does not belong to S , and that is replicated k times in $\text{Fan}(G, S, k)$. \square

As the following lemma shows, it is not very restrictive to consider only those fans in which the central subset S forms a clique. The advantage of restricting S in this way is that, when considering minors of the fan, we do not have to consider the ways in which such a minor might add edges between vertices of S .

Lemma 4. *Let G be a graph, let S be a subset of the vertices of G such that $G \setminus S$ is connected and every vertex in S has a neighbor in $G \setminus S$, and let G' be the graph formed from G by adding edges between every pair of vertices in S . Then there exists a constant c (depending on G and S) with the property that, for every $k > c$, $\text{Fan}(G', S, k - c)$ is a minor of $\text{Fan}(G, S, k)$.*

Proof. Let K be a maximum clique in the subgraph induced in G by S , and let $c = |S \setminus K|$. To form $\text{Fan}(G', S, k - c)$ as a minor of $\text{Fan}(G, S, k)$, simply contract one of the copies of G onto each vertex in $S \setminus K$. \square

For instance, in Figure 4(left), the three central vertices do not form a clique, but they can be made into a clique by contracting one of the six copies of the outer subgraph into the rightmost of the three central vertices. Thus, in this example, we may take $c = 1$.

4.2 Densest minors of fans

The following technical lemma will be used to compare the densities of minors of fans in which different copies of G are contracted differently from each other.

Lemma 5. *Let $a, b, c, d, e,$ and f be any six non-negative numbers, with $d + 2e, d + 2f,$ and $d + e + f$ positive. Then $(a + b + c)/(d + e + f) \leq \max\{(a + 2b)/(d + 2e), (a + 2c)/(d + 2f)\},$ with equality only in the case that the two terms in the maximum are equal.*

Proof. $(a + b + c)/(d + e + f)$ is a weighted average of the other two fractions with weights $(d + 2e)/(2d + 2e + 2f)$ and $(d + 2f)/(2d + 2e + 2f)$ respectively. As a weighted average with positive weights, it is not larger than the maximum of the two averaged values, and can be equal only when the two averaged values are equal. \square

With the assumption that S is a clique, any fan has a density-minimal minor that is also a fan with equal or greater density:

Lemma 6. *Let G be a connected graph, let S be a proper subset of G that induces a clique in $G,$ let every vertex in S be adjacent to at least one vertex in $G \setminus S,$ and let $G \setminus S$ induce a connected subgraph of $G.$ Then, for every $k,$ there exists a densest minor of $\text{Fan}(G, S, k)$ that is isomorphic to $\text{Fan}(G', S', k),$ where G' is some minor of G (that may depend on k) and S' is the image of S in $G'.$ In addition if S' is nonempty then G' can be chosen so that $\text{Fan}(G', S', k)$ is density-minimal; if S' is empty then G' can be chosen to be itself density-minimal.*

Proof. Let H be a densest minor of $\text{Fan}(G, S, k),$ and let S' be the image of S in $H.$ Note that it is not possible for all copies of G to be contracted onto S' in $H,$ for (with $s = |S'|$) the density $(s - 1)/2$ that would arise from this possibility is less than the density

$$\frac{s(s - 1)/2 + ks}{s + k} = \frac{s}{s + k} \cdot \frac{s - 1}{2} + \frac{k}{s + k}s$$

of the graph in which, in each copy of $G,$ the vertices that are not part of S are contracted into a single vertex: as the formula shows, this latter graph's density is a weighted average of $(s - 1)/2$ and $s,$ and therefore exceeds $(s - 1)/2.$

Suppose that the copies of G in the fan are transformed in H into two or more different minors, and let G_1 and G_2 be two of these minors. Define the integers $b, c,$ and a respectively to be the number of edges contributed to H by minors of type $G_1,$ the number of edges contributed to H by minors of type $G_2,$ and the number of remaining edges (including edges connecting vertices in $S'.$ Similarly, define $e, f,$ and d respectively to be the number of vertices contributed by minors of type $G_1,$ the number of vertices contributed by minors of type $G_2,$ and the remaining number of vertices (including vertices

in S'). Then the density of H is $(a + b + c)/(d + e + f)$, the density of the minor of $\text{Fan}(G, S, k)$ formed by replacing all copies of G_2 by G_1 is $(a + 2b)/(d + 2e)$, and the density of the minor formed by replacing all copies of G_1 by G_2 is $(a + 2c)/(d + 2f)$. By Lemma 5, one of these two replacements provides a minor of $\text{Fan}(G, S, k)$ that is at least as dense as H but that has one fewer different type of minor of G . By induction on the number of types of minors of G appearing in H , there is a densest minor of $\text{Fan}(G, S, k)$ of the form $\text{Fan}(G', S', k)$ where G' is a minor of G . Among all choices of G' leading to densest minors of this form, choose G' to be minimal in the minor ordering.

If $S' = \emptyset$, G' must be density-minimal, for any minor of G' with equal or smaller density would have been chosen in place of G' . If $S' \neq \emptyset$, $\text{Fan}(G', S', k)$ must be density-minimal, for (again using Lemma 5 to reduce the number of different types of minor in the graph) substituting some copies of G' for smaller minors could only produce a graph of equal density if these substitute minors could all be chosen to be isomorphic copies of a single minor G'' , but then (because of the choice of G' as giving the densest minor of this form) $\text{Fan}(G'', S'', k)$ would have equal density to $\text{Fan}(G', S', k)$, contradicting the choice of G' as a minor-minimal graph whose fan has this density. \square

5 Limiting density from density-minimal graphs

As we now show, the density-minimal graphs constructed in the previous two sections provide examples of minor-closed families, with the limiting density of the minor-closed family equal to the density of the density-minimal graph.

5.1 Minor-closed families with the density of a given density-minimal graph

For any graph G , define $\text{ComponentFamily}(G)$ to be the family of minors of graphs of the form $\text{Fan}(G, \emptyset, k)$ for some positive k . The graphs in $\text{ComponentFamily}(G)$ are characterized by the property that every connected component is a minor of G . Clearly, $\text{ComponentFamily}(G)$ is minor-closed.

Lemma 7. *If G is density-minimal, then the limiting density of $\text{ComponentFamily}(G)$ equals the density of G .*

Proof. $\text{ComponentFamily}(G)$ contains arbitrarily large graphs with this density, namely the graphs $\text{Fan}(G, \emptyset, k)$. Therefore its limiting density is at least equal to the density of G . But the fact that its limiting density is no more than the density of G is immediate, for if it contained any denser graph then that graph would have a dense component forming a minor of G and contradicting the assumed density-minimality of G . \square

Corollary 8. *The set of limiting densities of minor-closed graph families is a superset of the set of densities of density-minimal graphs.*

5.2 Gaps in the densities of density-minimal graphs

The idea of designing a minor-closed family to achieve a given density comes up in a different way in the following lemma, which is the basis for our proof that the limiting densities form a well-ordered set.

Lemma 9. *For every $\Delta \geq 0$ there exists a number $\delta > 0$ such that the open interval $(\Delta, \Delta + \delta)$ does not contain the density of any density-minimal graph.*

Proof. Let $\text{Sparse}(\Delta)$ be the minor-closed family of graphs that do not contain any minor that is more dense than Δ . By the Robertson–Seymour theorem, $\text{Sparse}(\Delta)$ can be characterized by a finite set of minimal forbidden minors; a minimal forbidden minor for $\text{Sparse}(\Delta)$ is a graph whose density is greater than Δ but all of whose proper minors have density less than or equal to Δ . Thus, any minimal forbidden minor of $\text{Sparse}(\Delta)$ must have density greater than Δ ; let $\Delta + \delta$ be the smallest density among these finitely many forbidden minors. Now suppose that G is a density-minimal graph whose density is greater than Δ . G is not in $\text{Sparse}(\Delta)$, so it has as a minor one of the forbidden minors of $\text{Sparse}(\Delta)$, all of which have density at least $\Delta + \delta$. Because G is density-minimal, its density is at least as great as that of its minor, and therefore is also at least $\Delta + \delta$. Therefore, G cannot have density within the interval $(\Delta, \Delta + \delta)$. \square

6 Separators in minor-closed graph families

We will use separators as a tool to find dense fans in any minor-closed family, so in this section we formalize the mathematics of separators in the form that we need.

If G is any graph, define a *separation* of G to be a collection of subgraphs that partition the edges of G . We call these subgraphs the *separation components* of a separation. Define the *separator* of a separation to be the set of vertices that belong to more than one separation component. The separator theorem of Alon, Seymour, and Thomas [1] can be rephrased in terms of separations:

Lemma 10 (Alon, Seymour, and Thomas [1]). *For any minor-closed graph family \mathcal{F} there exists a constant $\sigma_{\mathcal{F}}$ such that every n -vertex graph G in \mathcal{F} has a separation of G into two subgraphs G_1 and G_2 , each having at least $n/3$ vertices, with at most $\sigma_{\mathcal{F}}\sqrt{n}$ vertices in the separator.*

For our purposes we need a form of separator theorem that produces much smaller separation components, with only constant size. This can be achieved by repeatedly applying the separator theorem of Lemma 10 to larger components until they are small enough. A finer separation theorem of this type for planar graphs, proved in the same way from the planar separator theorem, has long been known [10, Theorem 3] and similar methods have been used as well in the context of minor-closed graphs (e.g. see [16, Lemma 3.4]). However, in order to control the density of the components, we need to be careful about how we measure the size of the separator, since a single separator vertex may appear in many components. We define the *separator multiplicity* of a vertex v in a separation of

G to be one less than the number of separation components containing v , so that a vertex has nonzero separator multiplicity if and only if it belongs to the separator, and we define the *total multiplicity* of a separation to be the sum of the separator multiplicities of the vertices. If a vertex has separator multiplicity zero we say that it is an *internal vertex* of its separation component.

Lemma 11. *For any minor closed graph family \mathcal{F} there exists a constant $\rho_{\mathcal{F}}$ such that, for every $0 < \epsilon < 1$ and every n -vertex graph G , there is a separation of G into subgraphs, each having at most $\rho_{\mathcal{F}}/\epsilon^2$ vertices, with total multiplicity at most ϵn .*

Proof. Set

$$X = \frac{2\sigma_{\mathcal{F}}}{\sqrt{2/3} + \sqrt{1/3} - 1} \text{ and } \rho_{\mathcal{F}} = 3X^2.$$

We prove by induction on n a stronger version of the lemma: whenever $n \geq (X/\epsilon)^2$, there exists a separation in which the size of each separation component is as specified and for which the total multiplicity is at most $\epsilon n - X\sqrt{n}$. For $n < (X/\epsilon)^2$ the induction does not go through because $\epsilon n - X\sqrt{n}$ is negative, but in this case the lemma itself follows trivially since we may choose a separation with a single separation component (all of G) and no vertices in the separator; the total multiplicity is 0, which is at most ϵn as desired. As a base case for the induction, whenever $(X/\epsilon)^2 \leq n \leq 3(X/\epsilon)^2 = \rho_{\mathcal{F}}/\epsilon^2$, we again choose a separation with a single separation component and no vertices in the separator; for n at least this large, $\epsilon n - X\sqrt{n} \geq 0$ which again exceeds the zero value of the total multiplicity.

For larger values of n , apply Lemma 10 to find two separation components G_1 and G_2 of G , with separator S . Each separation component will have at least $n/3 \geq (X/\epsilon)^2$ vertices, so we may apply the induction hypothesis to them, separating them into components with at most $\rho_{\mathcal{F}}/\epsilon^2$ vertices per component. Our separation of G is the union of the sets of components in these two separations.

The separator for this separation consists of the union of the separators in the two subgraphs G_1 and G_2 , together with the vertices in S . Let n_1 and n_2 be the numbers of vertices in G_1 and G_2 , respectively. When we combine the two separations, the separator multiplicity of vertices outside S remains unchanged, and the multiplicity of vertices in S increases by exactly one. Therefore, by the induction hypothesis, the combined total multiplicity is at most

$$(\epsilon n_1 - X\sqrt{n_1}) + (\epsilon n_2 - X\sqrt{n_2}) + \sigma_{\mathcal{F}}\sqrt{n}.$$

The worst case, for the terms $X\sqrt{n_1}$ and $X\sqrt{n_2}$ appearing in this expression, is that n_1 is $n/3$ and n_2 is only a small amount larger than $2n/3$; for, in that case, the sum of these two terms is as small as possible. If n_1 and n_2 are closer together, the sum of these terms is larger and a larger amount is subtracted from the total expression. Substituting n_1 and n_2 for this worst case and using the assumption that $\epsilon \leq 1$ to bound $\epsilon(n_1 + n_2) \leq \epsilon n + \sigma_{\mathcal{F}}\sqrt{n}$ allows this upper bound on the total multiplicity of the separator to be regrouped as

$$\epsilon n - (X\sqrt{2/3} + X\sqrt{1/3} - 2\sigma_{\mathcal{F}})\sqrt{n}.$$

Due to our choice of X , this further simplifies to $\epsilon n - X\sqrt{n}$ as required by the induction. \square

7 Density-minimal graphs from limiting density

As we now show, the limiting density of a minor-closed graph family \mathcal{F} may be approximated by the densities of a subfamily of the density-minimal graphs that it contains. Some care is needed, though, as it is possible for a minor-closed graph family to contain some density-minimal graphs whose densities are strictly larger than the limiting density of the family. The main idea of the sequence of lemmas in this section is to show that \mathcal{F} contains large dense graphs with a structure that is progressively more uniform, eventually leading to the result that it contains large dense density-minimal fans.

Lemma 12. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ and any integer $k > 0$, \mathcal{F} contains a graph G , such that G has a separation into k or more separation components, each of which has size $O(\epsilon^{-2})$ (where the constant hidden in the O -notation depends only on \mathcal{F}), such that the disjoint union of the separation components forms a graph with density at least $\Delta - \epsilon$.*

Proof. Let $\delta = \epsilon/(2\Delta)$. Because \mathcal{F} has limiting density Δ , it contains arbitrarily large graphs G of density at least $\Delta - \epsilon/2$. By Lemma 11, any such graph G with n vertices has a separation into separation components of size $O(\delta^{-2}) = O(\epsilon^{-2})$, in which the total multiplicity of the separation is at most δn . For sufficiently large n , the number of components is at least k . The density of the disjoint union of the separation components can be calculated by replacing the denominator, n , in the definition of the density of G by the new denominator $n + \delta n$; therefore, the density of the disjoint union is at most $(\Delta - \epsilon/2)/(1 + \delta) \geq \Delta - \epsilon$. \square

A simple argument related to Chebyshev's inequality lets us replace the density of the disjoint union of the separation components (a weighted average of their individual densities) by a lower bound on the density of each separation component.

Lemma 13. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ and any integer $k > 0$, \mathcal{F} contains a graph G , such that G has a separation into k separation components of size $O(\epsilon^{-2})$ in which each component has density at least $\Delta - \epsilon$. As in Lemma 12, the hidden constant in the O -notation depends only on \mathcal{F} .*

Proof. By Lemma 12, for any k' there is graph G in \mathcal{F} with a separation into k' separation components of size $O(\epsilon^{-2})$, such that the density of the disjoint union of the components is at least $\Delta - \epsilon/2$. But the density of the disjoint union is a weighted average of the densities of its separation components, weighted by the number of vertices in each component. The weights of any two components differ by at most a factor of $O(\epsilon^{-2})$. Therefore, if Δ^* denotes the largest density of any graph of size $O(\epsilon^{-2})$ in \mathcal{F} (necessarily bounded since there are only finitely many graphs of that size) then, in order to achieve a weighted average of $\Delta - \epsilon/2$, every $O(2\epsilon^{-3}(\Delta^* - \Delta + \epsilon/2))$ separation components with density

less than $\Delta - \epsilon$ (at least $\epsilon/2$ units below the average) must be balanced by at least one component with density greater than that threshold (and at most $\Delta^* - \Delta + \epsilon/2$ units above the average). By setting k' sufficiently large we may find a graph G' in which at least k of the separation components have density at least $\Delta - \epsilon$, and by deleting the lower-density components of G' we can find a graph with k dense separation components. As G' is a subgraph of G , it must belong to \mathcal{F} . \square

We define two separation components to be *isomorphic* if they are isomorphic as labeled graphs, with a labeling of each vertex according to whether it is an internal vertex of the component or a separator vertex.

Lemma 14. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ and any integer $k > 0$, \mathcal{F} contains a graph G , such that G has a separation into k isomorphic separation components of size $O(\epsilon^{-2})$ in which each component's density is at least $\Delta - \epsilon$. The hidden constant in the O -notation depends only on \mathcal{F} .*

Proof. Let $s = O(\epsilon^{-2})$ be the bound on the component size for \mathcal{F} and ϵ given by Lemma 13, and let N be the number of isomorphism classes of s -vertex labeled graphs in \mathcal{F} in which each vertex is labeled as an internal vertex or a separator vertex; note that, by Lemma 13, $N > 0$. By Lemma 13, there is a graph G' in \mathcal{F} with a separation into kN components of size at most s , all of which have density at least $\Delta - \epsilon$. At least k of these components must be isomorphic to each other; let G be the union of those isomorphic components, with all the other components removed. Then G has the stated properties, and as a subgraph of G' it belongs to \mathcal{F} . \square

The labeling used in Lemma 14 is not refined enough for our purposes. What we need is that, if v is a vertex of the separator of G , then every separation component uses v in the same way. More formally, we define a separation of a graph G to be *uniform* if every separation component in it is isomorphic to the same s -vertex graph (for some s), and if the vertices of G can be labeled with the integers from 1 to s in such a way that all components are isomorphic as labeled graphs. Equivalently, a separation is uniform if (for some $t \leq s$) the separation vertices may be given labels from 1 to t in such a way that the separation vertices within each separation component have distinct labels, and all pairs of separation components have an isomorphism that respects the labels. To achieve this, we use *color coding*, a variant of the probabilistic method developed for solving subgraph isomorphism problems [2].

Lemma 15. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ and any integer $k > 0$, \mathcal{F} contains a graph G , such that G has a uniform separation into k separation components of size $O(\epsilon^{-2})$ each of which has density at least $\Delta - \epsilon$. The hidden constant in the O -notation depends only on \mathcal{F} .*

Proof. Let $s = O(\epsilon^{-2})$ be the size of the components produced by Lemmas 13 and 14. By Lemma 14, there exists a graph G' with a separation into $k(s-1)^{s-1}$ isomorphic separation components of size at most s in which each component has density at least $\Delta - \epsilon$. Suppose that, for each separation component of G' , $t < s$ of the vertices are

separator vertices, and the remaining $s - t$ vertices are internal to the component. We then choose, independently and uniformly at random for each vertex of the separator of G' , a label from 1 to t , and we also choose (arbitrarily rather than randomly) a labeling L of a single isomorphic copy of the separation component (separate from G') that places distinct labels from 1 to t on its separator vertices. For each separation component of G' , the probability is at least t^{-t} that it has an isomorphism to L such that the labels of the component in G' match the labels in L , so the expected number of separation components with label-preserving isomorphisms of this type is at least $k(s - 1)^{s-1}t^{-t} \geq k$. Therefore, there exists a labeling of G' that matches or exceeds this expected value, and allows at least k of the separation components of G' to be given matching labels on their separator vertices. By keeping these k matching components and deleting the rest, we find a subgraph G of G' that belongs to \mathcal{F} and has a uniform separation as required. \square

Observe that, in a uniform separation of a simple graph, there can be no edges in which both endpoints belong to the separator, because that would result in the graph being a multigraph. Therefore, if the graph is connected, each separation component must have at least one non-separator vertex. In the labeling defining a uniform separation of a graph G , we say that the label of a separator vertex is *singular* if there is only one vertex in G with that label, and *plural* otherwise.

Lemma 16. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ and any integer $k > 0$, \mathcal{F} contains a graph G , such that G has a uniform separation into k separation components of size $O(\epsilon^{-2})$ in which each component has density at least $\Delta - \epsilon$ and in which each separator vertex label is singular. The hidden constant in the O -notation depends only on \mathcal{F} .*

Proof. Let $s = O(\epsilon^{-2})$ be the maximum size of the separation components produced by Lemma 15 for \mathcal{F} and ϵ . By Lemma 15, there exists a graph G' in \mathcal{F} with a uniform separation into k^{2^s} separation components of size at most s in which each component has density at least $\Delta - \epsilon$. We then consider each label of a separator vertex in G' in turn; for each such label, we either make it singular or make it not be a separator vertex any more, at the expense of reducing the number of separation components to the square root of its former value. Our choice of k^{2^s} separation components at the start of this process ensures that there will be at least k separation components when we have completed the process.

So, suppose that we have p^2 separation components in which label λ is plural, and we wish to find some subset of p of the components such that, for the graph induced by that subset, λ is either singular or not a separator vertex. We consider the number of different separator vertices that have the label λ . If this number is at least p , then we may choose a single separation component for each different vertex with that label, ending up with at least p separation components overall; the graph induced by these separation components remains uniformly separated and, in its separation, the vertices labeled λ are no longer separation vertices because each belongs to a single separation component. Making these vertices into internal vertices of the separation components does not change the density of the components.

In the other case, there are fewer than p separator vertices with label λ . Therefore, because there are p^2 separation components, one of the vertices labeled λ must be shared by at least p separation components. The graph induced by these separation components remains uniformly separated and in it λ is singular.

Repeating this refinement process once for each of the labels of separator vertices produces a uniform separation in which all separator vertex labels are singular, as required. \square

A graph with a uniform separation in which each separator vertex label is singular is just a more complicated way of describing a fan, and the density of the fan is at least as large as the density of the separation components it is formed from. However, in order to apply Lemma 6 we need the fan to have some more structure: the shared vertices of the fan should form a clique, the repeated parts of the fan should be connected, and each shared vertex should be adjacent to a repeated vertex.

Lemma 17. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ there exists a graph G , and a proper subset S of the vertices of G such that $\text{Fan}(G, S, k) \in \mathcal{F}$ for all positive k , such that $G \setminus S$ is connected and S forms a clique, such that each vertex of S is adjacent to a vertex in G , and such that for all sufficiently large k the density of $\text{Fan}(G, S, k)$ is at least $\Delta - \epsilon$.*

Proof. By Lemma 16 we can find, for any k , a graph G with $O(\epsilon^{-2})$ vertices and a proper subset S such that G has density at least $\Delta - \epsilon/2$ and $\text{Fan}(G, S, k)$ belongs to \mathcal{F} . By choosing G appropriately, we can additionally guarantee that $\text{Fan}(G, S, k')$ also belongs to \mathcal{F} for all $k' > k$; for, if each of the finitely many choices for G had a fixed bound on the number of times it could be repeated in a fan, this would contradict Lemma 16 for values of k that exceeded the maximum of these bounds. Because G and $\text{Fan}(G, S, k)$ come from a uniform separation, there are no edges in the subgraph induced by S . Let m be the number of edges in G , n be the number of vertices in G , s be the number of vertices in S , and m_i and n_i be the number of edges incident to and vertices in each connected component G_i of $G \setminus S$. Then $m/(n - s) \geq m/n \geq \Delta - \epsilon/2$ and $m/(n - s)$ is a weighted average of the quantities m_i/n_i , so if G_i is the component of $G \setminus S$ maximizing m_i/n_i then $m_i/n_i \geq \Delta - \epsilon/2$.

Let S_i be the set of vertices in S incident to G_i , and let G' have $G_i \cup S_i$ as its vertex set and have as its edges all the edges in G that are incident to G_i together with an edge between every two vertices in S_i . By Lemma 4, $\text{Fan}(G', S_i, k)$ belongs to \mathcal{F} for all k . As required, S_i forms a clique, $G' \setminus S_i$ is connected, and every vertex in S_i has a neighbor that is not in S_i . As k grows larger, the density of $\text{Fan}(G', S_i, k)$ converges to $m_i/n_i \geq \Delta - \epsilon/2$, so for all sufficiently large k it is larger than $\Delta - \epsilon$. \square

Dense fans in \mathcal{F} can be used to generate density-minimal graphs whose densities approximate the limiting density of \mathcal{F} .

Lemma 18. *Let \mathcal{F} be a minor-closed family with limiting density Δ . Then for any $\epsilon > 0$ there exists a density-minimal graph G in \mathcal{F} whose density belongs to the closed interval $[\Delta - \epsilon, \Delta]$.*

Proof. By Lemma 17 we can find G , S , and k_0 such that $\text{Fan}(G, S, k) \in \mathcal{F}$ for all positive k , such that $\text{Fan}(G, S, k)$ meets the conditions of Lemma 6, and such that the density of $\text{Fan}(G, S, k)$ is at least $\Delta - \epsilon$ for $k \geq k_0$. By Lemma 9 there exists δ such that there are no density-minimal graphs with densities in the open interval $(\Delta, \Delta + \delta)$. \mathcal{F} may contain graphs with density $\Delta + \delta$ or larger, but only finitely many of them; let k_1 be the number of vertices in the largest such graph, and set $k = \max(k_0, k_1 + 1)$. By Lemma 6 there exist G' and S' such that $\text{Fan}(G', S', k)$ is a densest minor of $\text{Fan}(G, S, k)$, and therefore also has density at least $\Delta - \epsilon$; since $\text{Fan}(G', S', k)$ has at least $k \geq k_1 + 1$ vertices (Lemma 3), it has density less than $\Delta + \delta$ and therefore its density is at most Δ . According to Lemma 6, G' and S' can be chosen so that either $\text{Fan}(G', S', k)$ is density-minimal, or G' is density-minimal and has the same density as the fan; in either case we have found a density-minimal graph in \mathcal{F} whose density lies in the desired range. \square

8 Main results

Theorem 19. *The set of limiting densities of minor-closed graph families is countable, well-ordered, and topologically closed.*

Proof. By the Robertson–Seymour theorem [13], according to which every minor-closed graph family can be characterized by a finite set of forbidden minors, there are countable many minor-closed families and therefore at most countably many limiting densities.

There can be no infinite descending sequence of limiting densities of minor-closed families, because if there were, there would be an infinite descending subsequence $\Delta_1, \Delta_2, \Delta_3, \dots$ that converged to some limit Δ^* . But then, for every δ , there would be a limiting density Δ_i within $\delta/2$ of Δ^* , and (by Lemma 18) a density-minimal graph G with density within $\delta/2$ of Δ_i , contradicting Lemma 9 according to which there is an open interval $(\Delta, \Delta + \delta)$ that does not contain the density of any density-minimal graph. Therefore, the set of limiting densities is well-ordered.

By well-ordering, any cluster point Δ of the set of limiting densities must be the limit of an increasing subsequence $\Delta_1 < \Delta_2 < \Delta_3 \dots$ of limiting densities of minor-closed families $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \dots$. But then $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3 \dots$ is also minor-closed and has Δ as its limiting density; hence the set of limiting densities contains all its cluster points and is topologically closed. \square

Theorem 20. *The set of limiting densities of minor-closed graph families is the topological closure of the set of densities of density-minimal graphs.*

Proof. By Corollary 8, the set of limiting densities of minor-closed graph families contains the set of densities of density-minimal graphs, and by Theorem 19 it contains the closure of this set. By Lemma 18 every limiting density of a minor-closed graph family belongs to the closure of the set of densities of density-minimal graphs. \square

Theorem 21. *Let Δ be the limiting density of a minor-closed graph family \mathcal{F} . Then $1 + \Delta$ is a cluster point in the set of limiting densities.*

Proof. Let $\text{Sparse}(\Delta)$ be defined as in the proof of Lemma 9 as the family of graphs with no minor denser than Δ . Define $\text{Apex}(\Delta, k)$ to be the family of graphs G with the following two properties:

- Each connected component G_i of G has at most k vertices, and
- Within each connected component G_i of G , one can find a vertex a_i (the *apex* of the component) such that $(G_i \setminus \{a_i\}) \in \text{Sparse}(\Delta)$.

$\text{Apex}(\Delta, k)$ is minor-closed, and as we show below, the sequence of graph families $\text{Apex}(\Delta, k)$ (for fixed Δ and variable k) has a sequence of limiting densities that converges to $1 + \Delta$ but that does not ever actually reach $1 + \Delta$.

Because the density of any graph in $\text{Sparse}(\Delta)$ is at most Δ , any connected component of a graph in $\text{Apex}(\Delta, k)$ that has i vertices has at most $\Delta(i-1)$ edges that are not incident to a_i , and another $i-1$ edges incident to a_i . Therefore, its density is at most $(1 + \Delta)\frac{i-1}{i}$, and the limiting density of $\text{Apex}(\Delta, k)$ is at most $(1 + \Delta)\frac{k-1}{k}$. In particular, it cannot be the case that any family $\text{Apex}(\Delta, k)$ has limiting density exactly equal to $1 + \Delta$.

For any ϵ there exists a density-minimal graph G in \mathcal{F} with density in the interval $[\Delta - \epsilon/2, \Delta]$, by Lemma 18. Because G is density-minimal and has density at most Δ , it belongs to $\text{Sparse}(\Delta)$. Let G have m edges and n vertices, and form a graph G' by connecting each vertex in G to a new vertex a . The graphs $\text{Fan}(G', \{a\}, k)$ have density $(m+n)k/(nk+1)$, which approaches $1 + m/n$ as k becomes large, so for some sufficiently large k the graph $H = \text{Fan}(G', \{a\}, k)$ has density at least $1 + \Delta - \epsilon$. The graphs $\text{Fan}(H, \emptyset, r)$ belong to $\text{Apex}(\Delta, nk+1)$: each component of these graphs has $nk+1$ vertices, meeting the limit on the component size in $\text{Apex}(\Delta, nk+1)$, and within each component of $\text{Fan}(H, \emptyset, r)$ the vertex a can be chosen as the apex. The family $\text{Apex}(\Delta, nk+1)$ contains arbitrarily large graphs $\text{Fan}(H, \emptyset, r)$ with density at least $1 + \Delta - \epsilon$, so its limiting density is within the interval $[1 + \Delta - \epsilon, (1 + \Delta)\frac{nk}{nk+1}]$ and the sequence of limiting densities of these families approaches $1 + \Delta$.

We have found a sequence of minor-closed families of graphs with limiting densities approaching but not equalling $1 + \Delta$, so $1 + \Delta$ is a cluster point in the set of limiting densities. \square

We say that a number Δ is an order-1 cluster point in the set of limiting densities if Δ is any cluster point of the set, and that Δ is an order- i cluster point if Δ is a cluster point of order- $(i-1)$ cluster points.

Theorem 22. *The subset of limiting densities of minor-closed graph families that are less than $3/2$ is the set*

$$\left\{ \frac{i}{i+1} \mid i \geq 0 \right\} \cup \left\{ \frac{3i+2j}{2i+j+1} \mid i \geq 1, j \in \{0, 1, 2\} \right\}.$$

For the integers $i \geq 1$ and $j \geq 1$, the numbers $i + (j-1)/j$ are order- i cluster points in the set of limiting densities.

Proof. The description of the limiting densities below $3/2$ follows from Theorem 20 and Lemma 2. The characterizations of the numbers $i + (j - 1)/j$ as order- i cluster points follows by induction on i using Theorem 21, with the description of the subset of limiting densities that are less than or equal to 1 as a base case for the induction. \square

9 Multigraphs

Instead of using simple graphs, the theory of graph may be formulated in terms of *multigraphs*, graphs in which two vertices may be connected by a *bond* of two or more edges and in which a single vertex may have any number of *self-loops*, edges that have only that vertex as its endpoints. In some ways, the theory for multigraphs is simpler: for instance, the forbidden minor for forests in the simple graph world is a triangle, whereas for forests it is a smaller graph with one vertex and one self-loop. This simplicity carries over to the theory of limiting densities as well.

Theorem 23. *The limiting density of a minor-closed family \mathcal{F} of multigraphs can only be an integer, a superparticular ratio $i/(i + 1)$ for a positive integer i , or unbounded.*

Proof. If \mathcal{F} includes bonds with arbitrarily large numbers of edges, then its limiting density is unbounded. Otherwise, there can be at most finitely many different multigraphs in \mathcal{F} that have a fixed number of vertices, the only ingredient needed to make the counting arguments in our proofs go through. So, multigraphs in \mathcal{F} obey a separator theorem (obtained by applying the Alon–Seymour–Thomas separator theorem to the underlying simple graph of the multigraph), they contain dense fans of multigraphs, and their densities are approximated by the densities of finite density-minimal multigraphs. But if a connected multigraph contains a cycle, self-loop, or bond, then its densest minor is just a single vertex with density equal to its rank, because in the multigraph world edge contractions do not lead to the removal of any other edges, and therefore can only increase the density of a graph with at least as many edges as vertices. Therefore, the only density-minimal multigraphs are the trees and the single-vertex multigraphs, and the only possible densities that can be obtained from them are the ones in the statement of the theorem. \square

More specifically, if \mathcal{F} can have a single connected component with arbitrarily high rank, its limiting density is unbounded. Otherwise, let r be the largest number such that graphs in \mathcal{F} can have an unbounded number of rank- r connected components. If r is nonzero, the limiting density of \mathcal{F} is r . If r is zero, and \mathcal{F} can have a single connected component with unbounded size, its limiting density is one. And if r is zero and all components of \mathcal{F} have bounded size, the limiting density of \mathcal{F} is $i/(i + 1)$, where i is the number of edges in the largest tree T such that $\text{Fan}(T, \emptyset, k) \in \mathcal{F}$ for all k .

10 Conclusions

We have investigated the structure of the set of limiting densities of minor-closed graph families; this set is topologically closed and well-ordered, contains cluster points of cluster points, and its exact members are known up to the threshold $3/2$. However, our results open many additional questions for investigation:

- Are all limiting densities rational?
- Is every limiting density equal to the density of a density-minimal graph, or are there cluster points in the set of limiting densities that are not themselves densities of density-minimal graphs? A positive answer to this question would also imply that all limiting densities are rational.
- If $x = p/q$ is a limiting density, where p and q are integers, can the size of the gap between x and the next larger limiting density be lower-bounded by a nonzero closed-form expression in terms of p and q ? Due to the existence of cluster points, some dependence on q seems to be necessary: a monotonic function of x alone would have to approach zero as x approaches the first cluster point, and could not give a nontrivial lower bound on the gaps for higher values of x .
- Are there any other cluster points between $3/2$ and 2 other than the ones of the form $2 - 1/i$ that we have already identified?
- What is the order type of the set of limiting densities? Since the limiting densities contain cluster points of any finite order, the order type is at least ω^ω ; is this the exact order type?
- What is the smallest cluster point of order i for each integer i ? Is it the number i itself?
- If Δ is a cluster point of limiting densities, must $\Delta - 1$ be a limiting density? If so, it would follow that all limiting densities are rational, that i is the smallest cluster point of order i , and that the order type of the set of limiting densities is ω^ω .

Acknowledgements

This work was supported in part by NSF grant 0830403 and by the Office of Naval Research under grant N00014-08-1-1015. We thank Éric Fusy for calling our attention to reference [3], and an anonymous referee for many helpful suggestions.

References

- [1] N. Alon, P. Seymour, and R. Thomas. A separator theorem for nonplanar graphs. *J. Amer. Math. Soc.* 3(4):801–808, 1990, doi:10.2307/1990903.

- [2] N. Alon, R. Yuster, and U. Zwick. Color-coding. *J. ACM* 42(4):844–856, 1995, doi:10.1145/210332.210337.
- [3] O. Bernardi, M. Noy, and D. Welsh. Growth constants of minor-closed classes of graphs. *J. Combinatorial Theory, Ser. B* 100(5):468–484, 2010, doi:10.1016/j.jctb.2010.03.001.
- [4] M. Chudnovsky, B. Reed, and P. Seymour. The edge-density for $K_{2,t}$ minors, *J. Combinatorial Theory, Ser. B*, to appear, doi:10.1016/j.jctb.2010.09.001.
- [5] R. Diestel. *Graph Theory*. Graduate Texts in Mathematics 173. Springer-Verlag, 2005, <http://diestel-graph-theory.com/>.
- [6] P. Erdős, A. Rényi, and V. T. Sós. On a problem of graph theory. *Studia Sci. Math. Hungar.* 1:215–235, 1966, http://www.renyi.hu/~p_erdos/1966-06.pdf.
- [7] A. V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica* 4:307–316, 1984, doi:10.1007/BF02579141.
- [8] A. V. Kostochka and N. Prince. On $K_{s,t}$ -minors in graphs with given average degree. *Discrete Math.* 308(19):4435–4445, 2008, doi:10.1016/j.disc.2007.08.041.
- [9] A. V. Kostochka and N. Prince. Dense graphs have $K_{3,t}$ minors. *Discrete Math.*, 2010, doi:10.1016/j.disc.2010.03.026.
- [10] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.* 9(3):615–627, 1980, doi:10.1137/0209046.
- [11] W. Mader. Homomorphieeigenschaften und mittlere Kantendichte von Graphen. *Math. Ann.* 174:265–268, 1967.
- [12] J. S. Myers and A. Thomason. The extremal function for noncomplete minors. *Combinatorica* 25(6):725–753, 2005, doi:10.1007/s00493-005-0044-0.
- [13] N. Robertson and P. D. Seymour. Graph Minors. XX. Wagner’s conjecture. *J. Combinatorial Theory, Ser. B* 92(2):325–357, 2004, doi:10.1016/j.jctb.2004.08.001.
- [14] E. Scheinerman and J. Zito. On the size of hereditary classes of graphs. *J. Combinatorial Theory, Ser. B* 61(1):16–39, 1994, doi:10.1006/jctb.1994.1027.
- [15] Z.-X. Song and R. Thomas. The extremal function for K_9 minors. *J. Combinatorial Theory, Ser. B* 96(2):240–252, 2006, doi:10.1016/j.jctb.2005.07.008.
- [16] S. Tazari and M. Müller-Hannemann. Shortest paths in linear time on minor-closed graph classes, with an application to Steiner tree approximation. *Discrete Applied Math.* 157(4):673–684, 2009, doi:10.1016/j.dam.2008.08.002.
- [17] R. Thomas and P. Wollan. The extremal function for 3-linked graphs. *J. Combinatorial Theory, Ser. B* 98(5):939–971, 2008, doi:10.1016/j.jctb.2007.11.008.
- [18] A. Thomason. An extremal function for contractions of graphs. *Math. Proc. Camb. Phil. Soc.* 95(2):261–265, 1984, doi:10.1017/S0305004100061521.
- [19] A. Thomason. The extremal function for complete minors. *J. Combinatorial Theory, Ser. B* 81(2):318–338, 2001, doi:10.1006/jctb.2000.2013.
- [20] H. Whitney. Non-separable and planar graphs. *Trans. Amer. Math. Soc.* 34(2):339–362, 1932, doi:10.2307/1989545.