

COVERING RECTILINEAR POLYGONS WITH AXIS-PARALLEL RECTANGLES*

V. S. ANIL KUMAR[†] AND H. RAMESH[‡]

Abstract. We give an $O(\sqrt{\log n})$ factor approximation algorithm for covering a rectilinear polygon with holes using axis-parallel rectangles. This is the first polynomial time approximation algorithm for this problem with an $o(\log n)$ approximation factor.

Key words. approximation algorithms, covering polygons

AMS subject classifications. 52C15, 68W25, 68W40

DOI. 10.1137/S0097539799358835

1. Introduction. We consider the problem of covering rectilinear polygons with axis-parallel rectangles. Given a rectilinear polygon P with complexity n (complexity refers to the minimum of the number of vertical edges and the number of horizontal edges in the polygon), this problem requires determining the minimum number of axis-parallel rectangles whose union covers P . The polygon P may have holes in it.

Applications. Cheng, Iyengar, and Kashyap [5] showed that this problem has applications to image compression. They claim that representing an image using a rectangle covering of its white pixels gives compression superior to that achieved by quadtrees. It also has applications to printing integrated circuits [9].

Hardness. Much effort has gone into determining the computational complexity of this problem. In spite of this, the exact complexity of this problem has remained open for many years and continues to do so. Masek [19] showed that this problem is NP-complete. Later, Culberson and Reckhow [6] used a clever reduction from 3-SAT to show that this is the case even when P has no holes. The next natural question is whether the number of rectangles needed to cover P can be computed *approximately*. Berman and Dasgupta [2] showed that this problem is MaxSNP-Hard for polygons with holes, ruling out the possibility of a polynomial time approximation scheme.

Approximation results. Note that the rectangle covering problem is a special case of the general set covering problem. Therefore, it admits an approximation algorithm with a performance guarantee of $O(\log n)$, using the greedy scheme due to Johnson [10] and Lovasz [16]. This was the best approximation factor known for the rectangle covering problem until now. Further, it is known that the general set covering problem cannot be approximated any better, modulo constant terms, unless $NP = P$ [17, 18]. However, this proof of hardness assumes certain properties about the set system which do not hold for the rectangle covering problem. In this paper, we address the issue of whether the $\Omega(\log n)$ approximation factor barrier can be broken in polynomial time for the rectangle covering problem.

There seem to be only a few examples of nontrivial algorithms breaking this

*Received by the editors July 9, 1999; accepted for publication (in revised form) March 31, 2003; published electronically October 2, 2003. An abstract of this work appeared in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, Atlanta, GA, 1999.

<http://www.siam.org/journals/sicomp/32-6/35883.html>

[†]CCS-5, Los Alamos National Laboratory, Los Alamos, NM 87544 (anil@lanl.gov). This work was done while this author was a graduate student at the Indian Institute of Science.

[‡]Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India 560012 (ramesh@csa.iisc.ernet.in).

barrier for specific instances of the set covering problem. Brönnimann and Goodrich [3] showed that the covering problem for any set system with Vapnik–Chervonenkis dimension d can be approximated within an $O(d \log(dc))$ factor, where c is the cost of the optimal covering. For the rectangle covering problem, while d is a constant, c can be shown to be $\Omega(\sqrt{n})$; therefore, the Brönnimann–Goodrich algorithm gives only an $O(\log n)$ approximation factor. Brönnimann and Goodrich [3] also showed that $O(c)$ sized set covers can be computed for two-dimensional disc covering and a problem related to three-dimensional polytope separation as these set systems admit ϵ -nets of small size. It is not clear whether the rectangle covering problem admits ϵ -nets of small size.

Special situations. There are special situations when the above barrier can indeed be broken. When P is hole-free, Franzblau [7] showed a factor 2 approximation guarantee. When P has holes, Franzblau also gave an $O(n \log n)$ time heuristic which gives an $O(\log n)$ approximation factor. When P is both vertically and horizontally convex (i.e., the intersection of any vertical or horizontal line and P is just a single line segment), Chaiken et al. [4] gave a polynomial time algorithm which computed the minimum number of rectangles required, exactly. This was improved upon by Franzblau and Kleitman [8], who achieved the same result under the weaker restriction that P is just vertically convex. Note that both restrictions preclude the presence of holes.

Other papers which have dealt with this problem are [9, 11, 12, 13, 5].

Our result. We give the first algorithm to break the $\Omega(\log n)$ barrier even when P has holes. Our algorithm gives an approximation guarantee of $O(\sqrt{\log n})$.

Our algorithm is in fact trivial, and our contribution lies entirely in showing a lower bound. For simplicity, assume that all holes in P are point holes.¹ Then our algorithm simply puts one rectangle for each *strip*, i.e., a stretch of points between two vertically aligned holes (e.g., strip s and its associated rectangle R in Figure 1; this is defined formally in section 2). This rectangle covers the strip entirely and is made as thick as possible. It is easy to see that the rectangles for all strips together cover P (the boundary of P must be treated as being lined by point holes for this). Also note that the rectangles for two distinct strips could be identical (e.g., strips B and D in Figure 2). We show that the total number of distinct rectangles $\#N$ obtained in the above process is $O(\sqrt{\log n} * |OPT|)$, where OPT is the minimum cover.

The lower bound. The main hurdle in breaking the $\Omega(\log n)$ barrier is to obtain a good lower bound for the optimum. One such lower bound is the cardinality of the largest *independent set* or *antirectangle*, i.e., a set of points in P , no two of which can be covered by the same rectangle. Chvatal (as reported in [4]) originally conjectured that the size of the minimum covering equals the size of the largest independent set. While this is indeed true for vertically and horizontally convex P , as shown by Chaiken et al. [4], it is not true for general P , with or without holes. Szemerédi found a counterexample with holes, and Chung found one without holes (both reported in [4]). Erdős (as reported in [4]) asked whether the ratio of the sizes of the minimum covering and the largest independent sets is bounded. It is easy to show that this ratio is $O(\log n)$. However, to the best of our knowledge, the best lower bound on this ratio known to date is just $21/17 - \epsilon$, due to [4].

Instead of using the above independent set bound, we use the *clique covering* lower bound. Consider the finite set of all points in P after suitable discretization.

¹The case when the holes are arbitrary can effectively be reduced to the case of point holes, as we will show later in this paper.

Consider the graph G with these points as vertices and an edge between two points if and only if they are both covered by some rectangle. It is easy to see that $|OPT|$ is exactly the size of the smallest clique cover of this graph (i.e., a collection of cliques that covers all vertices). We shall lower bound the clique cover number by obtaining an upper bound on the sizes of cliques in OPT .

One problem we face in the process is that the cliques of G could be very different in size and therefore do not admit a uniform upper bound. However, we show that if none of the strips are *jumpers* (we will define this term in section 3.2), then we can choose $\Theta(\#N)$ points such that the maximum clique size in the subgraph induced by them is $O(\sqrt{\log n})$, and therefore $|OPT| = \Omega(\#N/\sqrt{\log n})$, as required. Our proof of this fact is based on a somewhat detailed exploration of the structure of cliques in G .

There is a correspondence between the points that are chosen in the above induced subgraph and the set of strips. We choose two points for each strip, one to its left and one to its right, after partitioning the strips into disjoint sets called *families* (defined later); this ensures that the induced subgraph has $\Theta(\#N)$ points. As mentioned earlier, if there are no jumpers, we show that the maximum clique size is $O(\sqrt{\log n})$ (this is not strictly true; we show that a further subset of this set of points has this property). Also, the number of jumpers turns out to be at most $(|OPT|\sqrt{\log n})$, so we can ignore such strips and argue about the points defined by the rest. The above description is very incomplete and will be developed formally in the remaining sections.

A significant point to note is that while we show a lower bound on the clique covering number of graph G in the absence of jumpers, we are unable to show a good lower bound on the size of the largest independent set. We can show that the number of families is a lower bound on the size of the largest independent set; however, the average family size can be as large as $\Omega(\log n)$, as we shall show later.

Roadmap. In section 2, we show how to discretize the polygon and then describe our algorithm for laying rectangles. For simplicity and clarity, we first explain the arguments in section 3 under some restrictive assumptions. The general case requires a refinement of these ideas and is handled in section 4. Section 5 describes an example where the average family size is large. Section 6 mentions the loose threads which remain in this problem and also describes some related problems.

2. Preliminaries. Consider the grid formed by drawing infinitely long lines through each vertical and horizontal edge of the polygon (i.e., both the polygon boundary and the hole boundaries) (see Figure 1). Note that this need not be a uniform grid; the spacing between adjacent grid lines is not necessarily the same. Let n denote the *vertical complexity* of P , i.e., the number of horizontal grid lines. Without loss of generality, we assume that the vertical complexity is at most as large as the horizontal complexity.

Viewing the entire plane as partitioned into grid cells, the term *hole* shall henceforth denote any grid cell which is in the exterior of the polygon (i.e., either outside the outer boundary or within one of the holes). Since any grid cell lies either completely in the interior of the polygon or completely in the exterior of the polygon (see Figure 1), the above term is well defined.

Note that any two holes are either perfectly aligned or completely misaligned with respect to the grid lines. Also note that two holes could lie side by side, touching each other (e.g., holes A, B in Figure 1). If we treat each grid cell as a point, we get the case of point holes. However, the rest of the description will be in terms of cells.




FIG. 1. The grid: Dark cells and external cells are holes. C_1, C_2 form a strip s of length 2. R is the rectangle associated with s , and u and l are its upper and lower holes, respectively.

2.1. Defining strips and laying rectangles. A sequence of consecutive vertically aligned nonhole cells bounded by holes on the top and the bottom constitutes a *strip* (see Figure 1). The *length* of a strip A is the number of cells in it and is denoted by $l(A)$. The *upper hole* for a strip is the hole which lies immediately above the topmost cell for that strip. The *lower hole* is similarly defined.

For each strip, we define its *associated rectangle* to be the unique rectangle that covers this strip and extends as far as possible to the left and to the right. In other words, the associated rectangle is obtained by sweeping the strip to the left and right until it is blocked by some holes on both sides (see Figure 1). The algorithm simply adds the associated rectangle of every strip to the overall cover. Lemma 2.1 shows that these rectangles indeed cover the given polygon. The rest of the paper proves that this naive way of covering is at most an $O(\sqrt{\log n})$ factor larger than the optimum.

The hole which blocks the associated rectangle of a strip S on the right is called the *right blocking hole* of S , with ties broken in favor of the topmost hole. *Left blocking holes* are defined similarly.

LEMMA 2.1. *Each point in the polygon is contained in the associated rectangle of some strip.*

Proof. Each nonhole grid cell, c , belongs to a unique strip. This is because if the cell c is swept vertically up and down, it would hit a hole in both directions. The rectangle associated with this strip contains c . \square

Note that two strips could have identical associated rectangles (e.g., strips B, D in Figure 2). Consider equivalence classes of strips, where strips with the same associated rectangle are in one class. All but the rightmost of the strips in an equivalence class are called *unnecessary strips*. Clearly, an unnecessary strip can be ignored. It suffices to account for the rectangles associated with necessary strips.

2.2. Spanning, nestedness, and disjointness. A rectangle associated with strip A is said to *pass through* strip B if some but not all of the cells in strip B are contained in this rectangle (see Figure 2). In this situation, B is said to *l-span* A if it is to the left of A and *r-span* A if it is to the right of A . Note the strict condition in the above definition; i.e., the length of B must be strictly greater than that of A .




FIG. 2. A family: Strips A, B, C are in one right family and have the common blocking hole h . D is unnecessary and has the same associated rectangle as B . B r -spans A . Strips E and F are disjoint.




FIG. 3. A left clique. B l -spans A , and C l -spans B .

A collection of strips is called *right nested* if the strips increase in length from left to right and each strip r -spans all smaller strips. Left nesting is defined analogously. The strips A, B , and C in Figure 3 are left nested.

Two strips A, B are said to be *disjoint* if one of the following conditions holds: (i) the lower hole of A is horizontally aligned with or higher than the upper hole of B , or (ii) the upper hole of A is horizontally aligned with or below the lower hole of B (see Figure 2).

2.3. Successors, terminals, and witness cells. We bound the approximation factor of our algorithm by first identifying a subset of the nonhole cells and then showing lower bounds on the number of rectangles needed to cover these cells. The nonhole cells we identify are called *witness* cells. There are two kinds of witness cells, *left witness* cells and *right witness* cells. To define witness cells, we need the notion of successor strips.

We will formally define successor strips in section 3.1. Here, we introduce some properties of successors. The *right successor* strip for a strip A r -spans A , and the *left successor* for a strip A l -spans A . Not each necessary strip has a right successor; those strips which do not have right successors are called *right terminal* strips. *Left terminal* strips are defined analogously.

Given successors and terminals, witnesses are defined as follows. For each right terminal strip A , its right witness cell is the cell in A which is horizontally aligned with its right blocking hole (see Figure 4). For each right nonterminal strip A , its




FIG. 4. *Right witnesses for nonterminal strip A and terminal strip C .*

right witness is the cell in its right successor strip which is horizontally aligned with A 's upper hole (see Figure 4). The left witness is defined in a similar manner.

Note that right witnesses for right nonterminal strips A are not well defined if the upper hole of the right successor of A is aligned with the upper hole of A . An analogous fact holds for left witnesses of left nonterminal strips. This is clearly a problem because the above process may not define sufficiently many witness cells to obtain a good enough lower bound on the size of the optimum. We solve this problem by first identifying a constant fraction of the necessary strips with the following property: each such strip has a well-defined left witness. The remaining necessary strips are called *discarded* strips, and they play no role in the proof.

The precise definition of successor strips and the description of which strips are discarded appear in subsequent sections. We set up some more preliminaries in this section.

2.4. Cliques and the optimum cover. The optimum cover, denoted by OPT , must cover all the witness cells defined above. Consider a graph G whose vertices are the various witness cells defined above and whose edges denote that the two associated cells can be covered together by a single rectangle. Note that each cell is either completely inside or completely outside any maximal rectangle (one which has holes touching all four sides). Two witness cells are said to be *independent* if no single rectangle covers both of them; i.e., there is no edge between them in G . The following lemma holds.

LEMMA 2.2. *All witness cells contained in any rectangle form a clique in G . Conversely, any clique in G comprises witness cells which can be covered by just one rectangle.*

Proof. The first statement of the lemma follows from the definition of G .

We now show that any clique C in G can be covered by one rectangle. Let a, b, c, d be the leftmost, topmost, rightmost, and bottommost witness cells in C , respectively. Note that a, b, c, d need not be all distinct. We claim that the rectangle R having a, b, c, d on its left, top, right, and bottom edges, respectively, is hole-free (see Figure 5). The claim then follows since a, b, c, d are extreme points.

To show that R is hole-free consider all edges between witness cells a, b, c, d . If a, b, c, d are all distinct there will be six such edges, and fewer otherwise. There exist up to six rectangles, each of which covers both witness cells for one of the above edges (see Figure 5). Clearly, R is contained in the union of these rectangles. \square

We partition the witness cells into $O(|OPT|)$ cliques, where $|OPT|$ denotes the number of rectangles used by OPT , in the following manner. For each witness cell,




FIG. 5. The rectangle R formed by a, b, c, d and the six possible rectangles formed by each pair of witness cells, which together cover R .

assign it to some rectangle covering it in OPT , breaking ties arbitrarily. From Lemma 2.2 above, this corresponds to a partition of witness cells into cliques. We partition each clique further into two parts, one containing left witness cells and another containing right witness cells; these parts are called the *left clique* and the *right clique*, respectively.

Remark. We say that a strip is *in* a particular left clique (right clique, respectively) if its left witness cell (right witness cell, respectively) is in that clique, and, by abuse of notation, we will sometimes identify a strip with its witness cell.

To bound the approximation factor of our algorithm, we will show a lower bound on the number of cliques obtained above. This lower bound will exploit several interesting properties of these cliques. However, before we describe these properties, we need to specify how successors are determined and how strips to be discarded are identified. We start by describing the above for a simpler special case so as to bring out the intuition behind our proof.

We again remind the reader that unnecessary strips are being ignored, and any reference to a strip in the rest of the paper denotes a necessary strip.

3. The lower bound argument: A special case. We make the following assumptions in this section and illustrate the main ideas of the proof of the lower bound for this special case. We shall return to the general case in section 4.

Assumption 1. The length of each strip is a power of 2.

Assumption 2. The upper hole of a strip is not horizontally aligned with that of its left or right successor strip. The notion of a successor strip was introduced in section 2, and successors will be defined shortly.

Recall the preceding discussion on discarded strips in section 2. Assumption 2 precludes exactly those situations which forced us to introduce the notion of discarded strips. It follows that there is no need to discard any strips. Thus, all necessary strips will have associated left and right witnesses. Assumption 1 will make the definition of successor strips a little easier.

3.1. Defining families and successors. We organize strips into families as follows. Our proofs crucially exploit the interplay between cliques and families.

We define a *right family* to be a set of strips with the same right blocking hole (Figure 2). It is easy to see that strips in a right family are right nested (the strict increase in strip lengths from left to right is a consequence of the absence of unnecessary strips). The *right successor* of a strip A in a right family is defined as the next strip A' to the right in the family. The rightmost strip in the family does not have a successor and is a *right terminal strip*. Left families, successors, and terminal strips

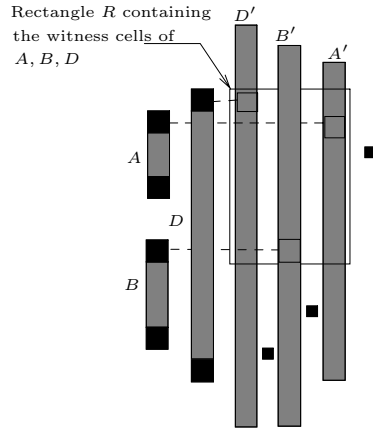


FIG. 6. A, B, D are in a right clique. A', B', D' are their respective right successors. A is an exception.

are defined analogously. It follows from Assumption 1 and the nestedness property that the number of strips in a family is at most $\log n$.

Remark. The notion of successor defined above will have to be changed when we deal with the general case in section 4 in the following manner. We will define two types of strips and use the above definition for strips of the first kind; strips of the second kind will require a different definition, and the successor of such a strip A will have length $< 2l(A)$. However, it will continue to be the case that the right successor of A r -spans A , and, similarly, the left successor of A l -spans A for all undiscarded strips A . Note that Assumption 1 precludes the existence of strips of the second kind above.

3.2. Properties of cliques. Our proof is based on some structural facts, which we state in the following lemmas. In section 3.3, we will use these lemmas to obtain the approximation factor. The proofs of these lemmas appear in section 3.4.

LEMMA 3.1. *All right (left, respectively) witness cells associated with right (left, respectively) terminal strips are independent. Therefore, the number of families and terminal strips is $O(|OPT|)$.*

All further references to strips in subsequent lemmas in this section will be to nonterminal strips.

LEMMA 3.2. *Strips in a right (left, respectively) clique belong to distinct right (left, respectively) families.*

LEMMA 3.3. *With the exception of at most one strip, all strips in a right (left, respectively) clique constitute a right (left, respectively) nested set of strips.*

Thus, each clique of OPT has at most two exceptions, one in each direction. The total number of exception strips is therefore $O(|OPT|)$. These exception strips can be removed from consideration. Figure 6 shows an example of an exception strip in a right clique. All further references to cliques in this section will assume that exception strips are not present.

LEMMA 3.4. *Strips in any right (left, respectively) clique are in distinct length categories; i.e., if a particular strip has length in the range $[2^i, 2^{i+1})$, then the next strip to the right (left, respectively) has length at least $2 \cdot 2^{i+1}$.*

²Strips in a clique will actually at least double in length by Assumption 1. However, we prefer

In addition, for any $x > 0$, the number of strips in a right (left, respectively) clique or a right (left, respectively) family whose length is at least 2^x times the length of one of the two previous strips to the left (right, respectively) is $O(\frac{\log n}{x})$.

LEMMA 3.5. Let A, B be strips in a particular right (left, respectively) clique, with $l(A) < l(B)$. Let A', B' be the right (left, respectively) successors of A, B , respectively. These four strips must be in the following order from left to right (right to left, respectively): A, B, B', A' . In addition, A' cannot r -span (l -span, respectively) B and must have its upper hole above that of B .

DEFINITIONS. A strip A is called a *right jumper* if its right successor has length at least $2^\Delta l(A)$, where Δ is a parameter. This parameter will be set to $\Theta(\sqrt{\log n})$ at the end. Left jumpers are defined analogously. The *vertical separation* between two holes a, b is the vertical distance between their lower boundaries, measured in terms of the number of grid cells (see Figure 12).

LEMMA 3.6. Let A, B be nonjumper strips in some right (left, respectively) clique, with $l(A) < l(B)$. The following two facts hold.

1. The vertical separation between the upper holes of A and B is at most $2^\Delta l(A)$.
2. If A is not amongst the smallest $\gamma\Delta$ nonjumper strips in this right clique (left clique, respectively), the vertical separation between the upper holes of A and B is at most $\frac{l(B)}{2^{(\gamma-1)\Delta}}$.

LEMMA 3.7. Let A and B be strips belonging to the same right (left, respectively) clique, with $l(A) < l(B)$. Then A lies completely above the right (left, respectively) blocking hole of B .

LEMMA 3.8. Let C and C' be the left and right cliques, respectively, containing the left and right witness cells, respectively, of strip A . Let B be a strip in C smaller than A . Let B' be a strip in C' smaller than A . Then B' cannot l -span B , and B cannot r -span B' .

3.3. Accounting for strips. The number of terminal strips is $O(|OPT|)$ by Lemma 3.1. The number of jumper strips is $O(|OPT| * \frac{\log n}{\Delta})$ by Lemma 3.4. All references to strips in the rest of this section are to nonterminal, nonjumper, nonexception (see Lemma 3.3) strips. All references to cliques assume that terminal, jumper, and exception strips have been removed; references to clique sizes denote sizes subsequent to this removal.

We now consider the remaining strips and show that there exists a large subset W of the witnesses associated with these strips such that the graph induced by W has only small, i.e., size $\Theta(\Delta)$, cliques. A rough reason why such a subset W exists is as follows.

By the nestedness property of strips in a clique and by Assumption 1, large cliques will necessarily have long strips and will therefore require proportionately large vertical space. In addition, Lemma 3.6 states that if A, B are strips in a large clique, then the vertical separation between the upper holes of A and B is small. Given the properties of cliques stated above, we will show that the left cliques containing A, B and the right cliques containing A, B cannot all satisfy the dual requirements of large vertical space and small vertical separation, unless one of these cliques is small. This intuition is formalized below.

Remark. We mention here that the rest of this section uses only Lemmas 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, and 3.8. Assumptions 1 and 2 will not be used directly. When we drop these assumptions and proceed to the general case, we will apply the

to work with this weaker condition, as it generalizes even when Assumption 1 is dropped.

description below to a carefully chosen subset of strips for which the above lemmas will indeed hold.

DEFINITION. The *right follower* of a strip B is defined to be the unique next longer strip, if any, in the right clique containing B (uniqueness holds from Lemma 3.3 and since exceptions are ignored). In Figure 6, D is a follower of B . Strip A also has D as its follower, but A is dropped as it is an exception strip. *Left followers* are defined similarly.

LEMMA 3.9. *Let P, Q be strips in some right clique C , with Q being the right follower of P (Q need not exist). Then one of the following must hold.*

1. P is among the smallest 3Δ strips in C , or P is the largest strip in C .
2. Q is among the smallest $5\Delta + 1$ strips in C'' , where C'' denotes the left clique containing Q .
3. $2^{\Delta}l(P) \leq l(Q)$.
4. Let C' denote the left clique containing P . Let P' be the strip in C' whose left follower is P . Either P' does not exist or $2^{\Delta}l(P') \leq l(P)$.
5. Let Q_1, \dots, Q_k , in increasing order of length, be the strips in C'' which are smaller than Q . Then $2^{2\Delta}l(Q_{k-1}) \leq l(Q)$.

Analogous statements hold for strips P, Q in a left clique C .

Proof. We suppose that none of the above five conditions holds to get a contradiction.

Since condition 1 is not satisfied, Q exists. Since condition 2 is not satisfied, $k \geq 5\Delta + 1$. By Lemma 3.8, none of Q_1, \dots, Q_k r -span P (see Figure 7(a)). Similarly, P cannot l -span any of Q_1, \dots, Q_k . By Lemma 3.3, Q r -spans P , and l -spans each of Q_1, \dots, Q_k ; therefore, the hatched regions must be hole-free. We consider two cases now, depending upon whether the upper hole of Q_k is above or below that of P .

First, suppose the upper hole of Q_k is aligned with or above that of P (see Figure 7(a)). The lower hole of Q_k must be above that of P ; otherwise, Q_k r -spans P , which contradicts Lemma 3.8. Therefore, the left blocking hole of Q_k will be horizontally aligned with or above the upper hole of P . By Lemma 3.7, Q_{k-1} is completely above this blocking hole and therefore completely above P . Since the fifth condition is not satisfied, $2^{2\Delta}l(Q_{k-1}) > l(Q)$. Then the vertical separation between P and Q is at least $l(Q_{k-1}) > \frac{l(Q)}{2^{2\Delta}}$. Since the first condition is not satisfied, Lemma 3.6 implies that the vertical separation between P and Q is at most $\frac{l(Q)}{2^{2\Delta}}$, a contradiction. While applying Lemma 3.6, recall that jumpers have been excluded from cliques earlier.

Second, suppose the upper hole of Q_k is below that of P (see Figure 7(b)). The bottom hole of Q_k must be below that of P ; otherwise, P will l -span Q_k , which contradicts Lemma 3.8. Since Q_k l -spans each of Q_1, \dots, Q_{k-1} , these must also have their upper holes below that of P . Since P cannot l -span Q_1, \dots, Q_k (by Lemma 3.8), their lower holes must also be below that of P . Since condition 4 is violated, P' exists. By Lemma 3.3, P l -spans P' , and therefore P' must be to the right of Q (see Figure 7(b)). Since P l -spans P' , Q l -spans each of Q_1, \dots, Q_k (by Lemma 3.3), and lower holes of Q_1, \dots, Q_k are below that of P , it must be the case that each of Q_1, \dots, Q_k either l -spans P' or is completely below it (Q_k l -spans P' , while Q_i lies below P' in Figure 7(b)). We claim that at most 2Δ of Q_1, \dots, Q_k can l -span P' . This is shown in the next paragraph. Then the vertical separation between the upper holes of $Q_{k-2\Delta}$ (which is completely below P') and Q is at least $l(P') > \frac{l(P)}{2\Delta} > \frac{l(Q)}{2^{2\Delta}}$ (because conditions 3 and 4 are not satisfied). Since $k - 2\Delta > 3\Delta$, Lemma 3.6 applied to $Q_{k-2\Delta}$ and Q implies that the vertical separation between the upper holes of these two strips is at most $\frac{l(Q)}{2^{2\Delta}}$, a contradiction.




FIG. 7. Two situations for the upper hole of Q_k .

It remains to show that at most 2Δ of Q_1, \dots, Q_k can l-span P' . We show that $Q_{k-2\Delta}$ cannot l-span P' . Suppose this is not true. Then $Q_{k-2\Delta}, \dots, Q_k$ all l-span P' and $l(Q_{k-2\Delta}) > l(P')$. By Lemma 3.4, $Q_{k-2\Delta}, \dots, Q_k$ are in distinct length categories, and therefore $l(Q) \geq 2^{2\Delta}l(Q_{k-2\Delta}) > 2^{2\Delta}l(P') > 2^\Delta l(P)$. The last inequality follows from the violation of condition 4. Then condition 3 is satisfied, a contradiction. \square

COROLLARY 3.10. *The number of nonterminal, nonjumper, nonexception strips is $O(|OPT| * (\frac{\log n}{\Delta} + \Delta))$.*

Proof. Each nonterminal, nonjumper, nonexception strip P must be in some right clique C and in some left clique C' .

We consider five classes of these strips, depending upon which of the conditions in Lemma 3.9 is satisfied. The number of strips P which satisfy the first condition is clearly $O(|OPT| * \Delta)$ because each clique in OPT has $O(\Delta)$ such strips. The number of strips P which satisfy the third condition is $O(|OPT| * (\frac{\log n}{\Delta}))$ by Lemma 3.4. Similarly, the number of strips P which satisfy the fourth condition is $O(|OPT| * \frac{\log n}{\Delta})$. Next, consider strips P which satisfy either condition 2 or 5. Such a strip P has a unique right follower Q . Note that any strip is the right follower of at most one strip. Thus it suffices to bound the number of strips Q which are right followers of strips P satisfying condition 2 or 5. Using the same argument as for condition 1, the number of strips Q satisfying condition 2 is $O(|OPT| * \Delta)$. Using an argument similar to that for condition 4, the number of strips Q satisfying condition 5 is $O(|OPT| * \frac{\log n}{\Delta})$. \square

Thus, given Assumptions 1 and 2, by setting $\Delta = \Theta(\sqrt{\log n})$ we get that the number of rectangles laid out by our algorithm is within an $O(\sqrt{\log n})$ factor of the optimal.

3.4. Proofs. We give the proofs of Lemmas 3.1, 3.2, 3.3, 3.4, 3.5, 3.6, 3.7, and 3.8 in that order. For the general case, we will not repeat the proofs of these lemmas. To convince the reader that these proofs would continue to hold even in the absence of Assumptions 1 and 2, we describe the proofs so that they are dependent only on the following facts and on proofs of previous lemmas in the above order, instead of Assumptions 1 and 2 directly. These facts are consequences of Assumptions 1 and 2 and of the way in which families and successors were defined. Thus, as long as the general case obeys these facts, and if we derive the generalizations of these lemmas in the same order, these proofs will continue to hold. However, there is one exception, namely the first part of Lemma 3.4, where we shall use Assumption 1. This shall




FIG. 8. Witnesses of terminal strips are independent.

require reproving when we get to the general case. This will be made precise in section 4.4.

We state Facts 1–6 here.

FACT 1. Strips in a right (left, respectively) family are right (left, respectively) nested and have the same right (left, respectively) blocking hole. Further, strips in distinct right (left, respectively) families have distinct right (left, respectively) blocking holes.

FACT 2. Each family has one terminal strip. Further, strips in a family lie in distinct length categories (length categories are given by the length ranges $[2^i, 2^{i+1})$, $1 \leq i \leq \log n - 1$).

Note that by Assumption 1, strips in a family satisfy a stronger property; namely, they at least double in size. However, the weaker property stated above will be all that is available in the general case.

FACT 3. The right (left, respectively) successor of a strip A is the next strip to the right (left, respectively) in the right (left, respectively) family containing A.

FACT 4. The right (left, respectively) successor of a strip A r-spans (l-spans, respectively) A.

FACT 5. The right (left, respectively) witness of a nonterminal strip A, if it exists, is horizontally aligned with the upper hole of A and lies on the right (left, respectively) successor of A.

Note that the condition “if it exists” always holds by Assumption 2. However, this will not be true in the general case, after Assumptions 1, 2 are dropped.

FACT 6. The right (left, respectively) witness of a terminal strip lies on the terminal strip itself and is horizontally aligned with its right (left, respectively) blocking hole.

We now give the proofs. At the end of each proof below, we make a careful note of which of the above facts are used.

Proof of Lemma 3.1. We give the proof for right terminals. The proof for left terminals is similar.

Suppose for a contradiction that there are two right terminal strips P, Q whose right witnesses are not independent. P and Q must be in distinct right families, and their right blocking holes are distinct, by Facts 1 and 2. Recall that the right witness cells of P and Q are in P and Q , respectively, and are horizontally aligned with their respective right blocking holes (see Fact 6).

Consider a rectangle containing the two witness cells; such a rectangle exists by Lemma 2.2 (see Figure 8). Clearly, both P and Q must have their upper holes above and lower holes below this rectangle. Without loss of generality, assume P is to the left of Q . Note that the right blocking holes for both P and Q must be in the hatched




FIG. 9. Two situations for M, N in Lemma 3.3 if N neither r -spans M nor is disjoint from M .

region (i.e., the extension of the above rectangle to the right; the rectangle itself must be hole-free). Therefore, both blocking holes must be to the right of Q . Then it follows that whichever one of these holes occurs further to the right of the other cannot be the right blocking hole for either P or Q . In addition, if the two blocking holes are vertically aligned, the lower one cannot be the right blocking hole for either P or Q .³ This gives a contradiction. \square

Remark. Note that the above proof used only Facts 1, 2, and 6 among the six facts listed above.

Proof of Lemma 3.2. We show that if A and B are two strips in the same right family, their right witness cells are independent. Similar arguments hold for the right.

Without loss of generality, assume that $l(A) < l(B)$. By Fact 1, B is to the right of A and r -spans A . By Facts 3 and 5, A defines its right witness cell either on B (if B is the right successor of A) or to the left of B . Since B defines its right witness cell on the horizontal line containing the upper hole of B and to the right of the upper hole of B , it is easy to see that any rectangle containing the right witness cells of A and B has to contain the upper hole of B . This implies that these two witness cells are independent and cannot be in a clique. \square

Remark. Note that the above proof used only Facts 1, 3, and 5 among the six facts listed above.

Proof of Lemma 3.3. We prove the lemma for a right clique C ; similar arguments hold for left cliques. The following fact will be useful.

FACT. Suppose strips M and N are two strips in C , and M is either to the left of N or vertically aligned with it. Then either M and N are disjoint or N r -spans M . This must be true; otherwise, one of the two situations shown in Figure 9 holds, and then M 's right witness cell will be independent from the right witness cell of N (recall Facts 4 and 5).

Suppose the strips in C are not right nested. By the above fact, if each strip in C r -spans the smallest strip in C , then there are no disjoint strips in C , and the strips in C must be right nested. So there must exist a strip in C which is disjoint from the smallest strip A in C ; consider the smallest such strip B . Clearly, neither A nor B can r -span any strip in C . We will show that all other strips D in C must r -span the lower of A, B . It would then follow from the above fact that the strips in C with the upper of A, B removed are right nested. We consider the case when B is below A ; the other case is similar.

Consider a strip D in C , other than the strips A, B , and suppose for a contradiction that D does not r -span B . Recall from the previous paragraph that neither

³Here, we use the fact that ties for the blocking hole were broken in favor of the upper hole.

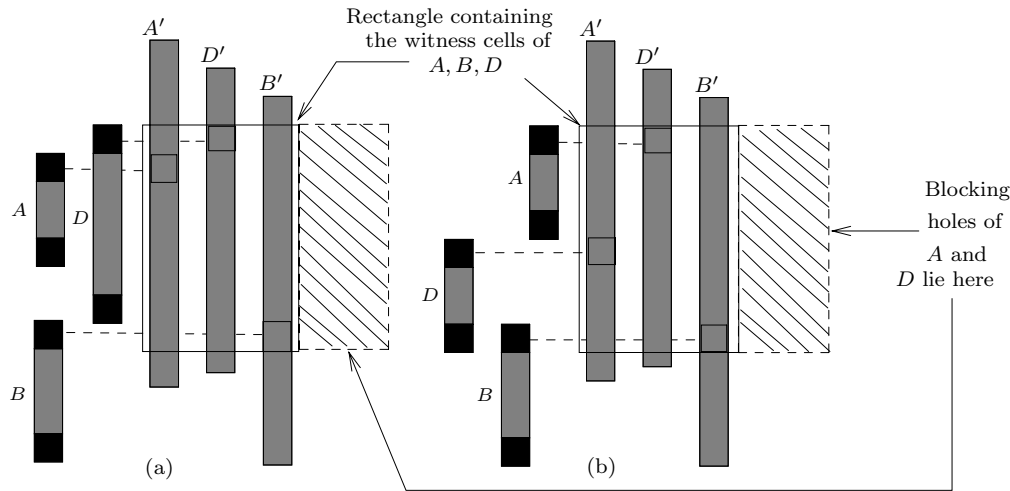


FIG. 10. (a) D r -spans A but not B . (b) A, B, D are disjoint. The order in which the successors A', B', D' are present is not important.

A nor B r -span any other strips in C . By the above fact, D must be disjoint from B . Further, D either is disjoint from A or is to the right of A and r -spans A (see Figure 10(a),(b)). We will show in the next paragraphs that the right blocking holes of the two upper strips among A, B, D must be identical. But, by Lemma 3.2, A, B, D must all be in distinct right families, and therefore, by Fact 1, they have distinct right blocking holes, a contradiction.

To show that the right blocking holes of the two upper strips among A, B, D must be identical, consider the rectangle R associated with clique C (by Lemma 2.2, such a rectangle exists). This rectangle has the following properties. The right witness cells for A, B, D are all in R , and A, B, D are themselves to the left of R . By Fact 5, the right witness cells of A, B, D are aligned with their respective upper holes and lie on their respective right successors, A', B', D' . By Fact 4, A', B', D' r -span A, B, D , respectively. Therefore, R must have its upper edge above the upper holes of A, B, D and its lower edge below all these holes. In addition, A', B', D' must all have their upper holes above R and bottom holes below R (see Figure 10); these successor strips must stab vertically through R . Note that the relative placement of A', B', D' is not important, though Figure 10 shows D' placed between A' and B' . Recall again that A, B are disjoint, D, B are disjoint, and D either r -spans A or is disjoint from A .

It follows from these properties that the right blocking holes of the upper two strips X, Y amongst strips A, B, D must be in the hatched region, i.e., the right extension of R . Next, since a strip and its right successor must have the same right blocking hole by Facts 1 and 3, the leftmost hole in the hatched region must be the right blocking hole for both X, Y , as required. \square

Remark. Note that the above proof directly used only Facts 1, 3, 4, and 5 among the six facts listed above, and Lemma 3.2.

Proof of Lemma 3.4. Strips in a right (or left) clique strictly increase in length by Lemma 3.3. The first part of the lemma then follows from Assumption 1.

Next, we prove the second part of Lemma 3.4. Right families are right nested by Fact 1. By Lemma 3.3, right cliques are also right nested. Therefore, in both right




FIG. 11. (a) B and B' lie between A and A' . (b) A' lies to the left of B . (c) A' r -spans B . (d) Upper hole of A' lies below that of B .

families and right cliques, strip lengths increase strictly monotonically to the right. Since the smallest strip length is 2^0 and the largest is $2^{\log n}$, the lemma follows. A similar argument holds for left families and cliques. \square

Remark. Note that the above proof directly used Assumption 1, only Fact 1 among the six facts listed above, and Lemma 3.3.

Proof of Lemma 3.5. We prove the lemma for a right clique C . The argument for left cliques is similar.

From Lemma 3.3, it follows that B r -spans A and hence must lie to the right of A . If B is to the right of A' , then the right witness cells a and b of A and B , respectively, are independent (see Figure 11(b); also see Facts 4 and 5). Therefore, B is to the left of A' and to the right of A . To show that the right successor of B also lies between A and A' , we show in the next paragraph that A' cannot r -span B . Since A' and B both r -span A , they cannot be disjoint either. Then it follows that the right blocking hole of B is vertically aligned with or to the left of A' . Therefore, the right successor of B is also to the left of A' (since, by Facts 1 and 3, a strip and its right successor have the same right blocking hole).

Suppose A' r -spans B . Then the right blocking hole of B would be the same as that of A (which is identical to that of A' by Facts 1 and 3) (see Figure 11(c); the hatched region must be hole-free); then A and B would be in the same right family by Fact 1. This contradicts Lemma 3.2.

It remains to show that the upper hole of A' is above that of B . By Fact 5, B 's right witness cell lies on the horizontal line containing the upper hole of B . If the upper hole of A' is horizontally aligned with or below that of B , the witness cells a and b of A and B , respectively, would be independent (see Figure 11(d)). Hence, the upper hole of A' lies above that of B . \square




FIG. 12. Vertical separation in a right clique.

Remark. Note that the above proof directly used only Facts 1, 3, 4, and 5 among the six facts listed above, and Lemmas 3.2 and 3.3.

Proof of Lemma 3.6. We prove the lemma for a right clique C . An analogous argument holds for left cliques.

First, consider part 1. Let A' be the right successor of A . By Fact 4, A' r -spans A . By Lemma 3.3, B must r -span A . By Lemma 3.5, the upper hole of A' is above that of B . It follows that the vertical separation between the upper holes of B and A is at most $l(A')$ and $l(A') \leq l(A)2^\Delta$, because A is not a right jumper. Part 1 follows.

Next, consider part 2. Let D be the smallest nonjumper strip in C , and let D' be the right successor of D (see Figure 12). D' exists because D is not a terminal. We will show in the next paragraph that the vertical separation between the upper holes of B and A is at most $l(D')$. Since D is not a right jumper, $l(D') \leq l(D)2^\Delta$. Further, since there are at least $\gamma\Delta$ strips smaller than A in C , using the increase in length categories given by Lemma 3.4, we get $l(D') \leq l(D)2^\Delta \leq \frac{l(B)}{2^{(\gamma)\Delta}}2^\Delta = \frac{l(B)}{2^{(\gamma-1)\Delta}}$, as required.

It remains to show that the vertical separation between the upper holes of B and A is at most $l(D')$. By Lemma 3.3, A and B both r -span D . By Lemma 3.5, A and B are between D and D' . Since D' is the right successor of D , D' must r -span D (see Fact 4); therefore, the lower hole of D' is below the upper hole of A . Further, by Lemma 3.5, the upper hole of D' is above that of B . It follows that the vertical separation between the upper holes of B and A is at most $l(D')$. \square

Remark. Note that the above proof directly used only Fact 4 among the six facts listed above, and Lemmas 3.3, 3.4, and 3.5.

Proof of Lemma 3.7. We prove the lemma for a right clique C ; the argument for the left is analogous.

By Lemma 3.3, B r -spans A . Let A' be the right successor of A and B' that of B . By Fact 4, A' r -spans A . By Lemma 3.5, B and B' are between A and A' , A' cannot r -span B , and the upper hole of A' is above that of B . It follows that the right blocking hole d of B must be to the left of A' (see Figure 13) and d can be only in the two hatched regions in the figure. But if it is in the upper of these two regions, by Fact 5, the right witness cells for A and B are independent. Therefore, it must be in the lower hatched region, which is completely below A . \square

Remark. Note that the above proof directly uses only Facts 4 and 5 among the six facts listed above, and Lemmas 3.3 and 3.5.




FIG. 13. Possible locations of the blocking hole of B .

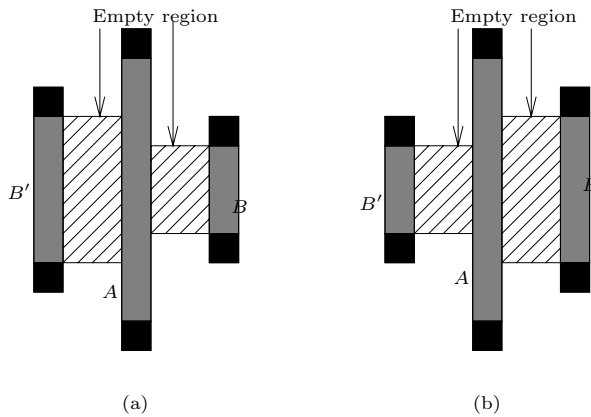


FIG. 14. Possibilities precluded by Lemma 3.8: (a) B' l -spans B . (b) B r -spans B' .

Proof of Lemma 3.8. Suppose B' l -spans B (Figure 14(a)). The hatched region must be hole-free since A l -spans B and r -spans B' (see Lemma 3.3). This means that B has its left blocking hole to the left of B' . Next, B 's left successor is either to the left of B' or to the right of A . This is true because any strip between B' and A which l -spans B must r -span B' as well and cannot have its left blocking hole to the left of B' (recall from Facts 1 and 3 that B and its left successor must have the same left blocking hole). Using Facts 4 and 5, it follows that the left witness cell of B is independent from that of A , a contradiction. Therefore B' cannot l -span B .

By an argument symmetric (see Figure 14(b)) to the one above, B cannot r -span B' . \square

Remark. Note that the above proof directly used only Facts 1, 3, 4, and 5 among the six facts listed above, and Lemma 3.3.

4. The general case: Removing ill-behaved strips. We now need to handle the general case. Assumption 2 is not very hard to handle; it is possible to show that a good fraction of strips define witnesses on at least one of the two sides. Assumption 1 is the most severe: strips of arbitrary lengths could result in large families and cliques in which the progression of strip lengths described by Lemma 3.4 is absent. This




FIG. 15. *The two kinds of strips considered in Lemma 4.1.*

forces us to alter the notion of successor/family defined in this section and partition strips into classes based on the nature of left and right successors. A detailed analysis is then needed for each of these classes.

Recall that we have already identified and decided to ignore unnecessary strips. The aim now is to discard some more ill-behaved strips and then redefine successors and families for the strips which remain. These remaining strips will constitute a constant fraction of the number of the rectangles laid out by our algorithm, and each such strip will have the property that it defines either a left witness cell or a right witness cell or both. Unlike the previous special case, it will no longer be true that each strip is part of some family. Each strip will still have a successor, though.

4.1. Discarding strips. Consider two sets of strips. The first set comprises necessary strips A with the property that A_l exists, where A_l is the closest necessary strip to the left, if any, such that the upper holes of A_l and A are horizontally aligned (Figure 15(a)) and A_l l-spans A . The second set comprises necessary strips A with the property that A_r exists, where A_r is the closest necessary strip to the right, if any, such that the upper hole of A_r and A are horizontally aligned (Figure 15(b)) and A_r r-spans A . By Lemma 4.1 below, the strips in the smaller of the above two sets can be ignored. Without loss of generality, assume that the former set is smaller. Let \mathcal{S}' denote the set comprising the remaining strips. For all strips A in \mathcal{S}' , A_l does not exist. This property, along with the subsequent definition of successors, will ensure that left witnesses are always defined for all strips in \mathcal{S}' .

LEMMA 4.1. *One of the above two sets must have size at most half the number of necessary strips.*

Proof. This follows because if A is in the first set then A_l is not in the second. \square

Defining categories, doubling strips and nondoubling strips. We classify strips in \mathcal{S}' into categories based on length. All strips with length in the range $[2^i, 2^{i+1})$ are in the i th category, $i \leq \log n - 1$.

Consider a strip A in \mathcal{S}' . Let B be the closest strip to the right of A in \mathcal{S}' , if any, which r-spans A and is in the same category as A . If B exists, then A is said to be *right nondoubling*, and B is said to be the *right successor* of A . In this case, any strip in \mathcal{S}' to the right of A and to the left of B which r-spans A must be in a higher category than A and therefore has its right blocking hole vertically aligned with or to the left of B . All other strips A are called *right doubling* strips. Analogous notions are defined to the left. As will be shown in Lemma 4.5 shortly, each strip is the right successor of at most one right nondoubling strip.

Note that we have not yet defined successors for right/left doubling strips. We will do so after we redefine families later in this section.

Remark. The strips considered in sections 3 and 3.1 are right and left doubling because of Assumption 1. As remarked in section 3.1, the notion of successor defined there is different from that defined above. The successor defined earlier corresponds to the successor of a right or left doubling strip.

LEMMA 4.2. *The number of strips in \mathcal{S}' which are both right nondoubling and left nondoubling is at most $|\mathcal{S}'|/2$.*

Proof. Consider strip A , which is both right nondoubling and left nondoubling. Let B be its right successor. Note that no other right nondoubling strip has B as a right successor (this will be shown formally in Lemma 4.5). We show in the next paragraph that B cannot be a left nondoubling strip. It follows that for each A which is both right nondoubling and left nondoubling, there is a unique B which is left doubling. The lemma follows.

Suppose B is left nondoubling. Let D denote its left successor. Then D is between A and B and l -spans B . Since B r -spans A , D r -spans A as well. Since A, B, D are all in the same category, D would be the right successor of A , a contradiction. \square

Let \mathcal{S} denote the subset of \mathcal{S}' comprising strips which are either left doubling or right doubling or both. The following lemma shows that \mathcal{S} contains at least one-fourth of all the necessary strips. All further references to strips in the paper will be to the strips in \mathcal{S} . We will account for only these strips; the remaining strips are discarded.

LEMMA 4.3. *The number of strips in \mathcal{S} is at least one-fourth the number of rectangles laid out to cover the polygon.*

Proof. Three kinds of strips have been ignored so far in defining the set \mathcal{S} :

1. unnecessary strips;
2. strips A with the property that A_l exists, where A_l is the closest necessary strip to the left such that the upper holes of A_l and A are horizontally aligned (Figure 15(a)) and A_l l -spans A ;
3. strips A , which are both left and right nondoubling.

Unnecessary strips have the same associated rectangle as some necessary strip. By Lemma 4.1, the number of strips of the second type is at most half the total number of necessary strips. Finally, by Lemma 4.2, the number of strips of the last type is at most half of the remainder obtained by removing strips of the first two types. \square

4.2. Defining families and witnesses. A crucial difference from before is that right families comprise only right doubling strips, and similarly for left families.

A *right family* is defined to be a set of right doubling strips in \mathcal{S} with the same right blocking hole. The right successor of a right doubling strip is the next strip A' to the right in the family. The rightmost strip in a right family has no successor, and is called a right terminal strip, as before. Left families and successors are defined similarly. Notice that this is the same definition as in section 3.1, where all strips were left and right doubling.

We have now defined successors for all strips, whether doubling or not. As before, the right (left, respectively) witness cell of a nonterminal strip A (whether doubling or nondoubling) is the cell in its right (left, respectively) successor, horizontally aligned with the upper hole of A . This witness is not defined if the upper holes of A and its successor are horizontally aligned. The right (left, respectively) witness cell of a right (left, respectively) terminal strip A is defined as before; i.e., it is the cell in A which is horizontally aligned with its right (left, respectively) blocking hole.

The next lemma shows that each strip in \mathcal{S} defines a left witness and is a straightforward consequence of the definition of \mathcal{S}' .




FIG. 16. Two possible configurations of strips A, B, C , where B is the right nondoubling successor of both A and C .

LEMMA 4.4. For all nonterminal strips $A \in \mathcal{S}$, the upper hole of A is not horizontally aligned with the upper hole of its left successor. Therefore, each strip A defines a left witness cell and possibly a right witness cell.

LEMMA 4.5. Any strip is the right (left, respectively) successor of at most two strips, one right (left, respectively) nondoubling and one right (left, respectively) doubling.

Proof. We prove the “right” case; the “left” case is analogous. First, suppose there are two right nondoubling strips A, C whose right successor is strip B . Then A, B, C belong to the same category, and hence $l(A) < l(B) < 2l(A)$ and $l(C) < l(B) < 2l(C)$. B must r-span both A and C . This can happen only either if A r-spans C (or vice versa) or A and C are both disjoint (see Figure 16). The former possibility cannot arise, since in that case the right successor of C would be A (or vice versa), and in the latter case, $l(B) \geq l(A) + l(C) \geq 2 \min\{l(A), l(C)\}$, which implies that B cannot be in the same category as either A or C , a contradiction.

Next, consider right doubling strips. If the right successor of a right doubling strip A is strip B , then A and B must belong to the same right family and B is the first strip to the right of A in its family. So B cannot be the right successor of any other right doubling strip. \square

4.3. Classifying strips. Based on the above, we can classify strips in \mathcal{S} as follows. Recall that strips in \mathcal{S} are left doubling or right doubling or both.

Class 1 This class contains all strips which are either terminal strips or jumper strips.

This class has two subclasses.

Class 1.1 This class contains all strips which are either right terminal strips or left terminal strips.

Class 1.2 This class contains all strips which are either right jumper strips or left jumper strips.

Class 2 This class includes all strips not in Class 1 and which define witness cells in both directions. This class has three subclasses.

Class 2.1 This class contains strips which are left doubling and right doubling.

Class 2.2 This class contains strips which are left doubling and right nondoubling.

Class 2.3 This class contains strips which are left nondoubling and right doubling.

Class 3 This class includes all strips not in Class 1 and which define only left witness cells. This class has two subclasses.

Class 3.1 This class contains strips which are left doubling.

Class 3.2 This class contains strips which are left nondoubling.

The aim now is to bound the number of strips in each class by $O(|OPT|(\Delta + \frac{\log n}{\Delta}))$.

4.4. Properties of cliques: The general case. This section deals with the properties of cliques in the general case. We will generalize all the lemmas stated in section 3.2. We partition the optimal clique cover, OPT , into disjoint cliques in the same way as described in section 2.4.

DEFINITIONS. Define $T_r(C, i)$ to be the set of strips of class i , i being one of 1.1, 1.2, 2.1, 2.2, 2.3, 3.1 or 3.2, whose right witness cells are in clique C of OPT . $T_l(C, i)$ for left witness cells is defined similarly. As mentioned earlier, we will sometimes abuse notation and identify a strip with its witness cell.

Before continuing, we reiterate that all right families comprise right doubling strips only, and, similarly, left families comprise left doubling strips only. Note that with our new definition of successors and families for strips in S , left versions (right versions, respectively) of Facts 1–6 of section 3.4 continue to hold for left doubling (right doubling, respectively) strips. This will allow us to use the proofs in section 3.4 here as well for these strips. In addition, it is clear that Facts 4 and 5 continue to hold for all cliques (and not just cliques associated with doubling strips). Therefore, the lemmas in section 3.2 which use only these facts will generalize to both the doubling and the nondoubling cases directly.

The following lemma is exactly Lemma 3.1 stated for the new definition of terminal strips. The proof is exactly the same as Facts 1, 2, and 6 continue to hold for doubling strips.

LEMMA 4.6. *All right (left, respectively) witness cells associated with right (left, respectively) terminal strips are independent. Therefore, the number of terminal strips is $O(|OPT|)$.*

The following lemma is exactly Lemma 3.2, stated in terms of the classes defined above. Again, the proof is exactly the same, because Facts 1, 3, and 5 continue to hold for doubling strips.

LEMMA 4.7. *All strips in $T_r(C, i)$ ($T_l(C, i)$, respectively) belong to distinct right families (left families, respectively) for all cliques C in OPT and i being one of 2.1, 2.3 (2.1, 2.2, 3.1, respectively).*

The following lemma, a generalization of Lemma 3.3, proves nested structure for cliques (modulo exceptions).

LEMMA 4.8. *For any clique C in OPT and any class i , i being one of 2.1, 2.2, 2.3 (2.1, 2.2, 2.3, 3.1, 3.2, respectively), there exists at most one strip (called the exception) whose removal makes $T_r(C, i)$ ($T_l(C, i)$, respectively) right nested (left nested, respectively).*

Proof. The proof for $T_l(C, i)$, i being one of 2.1, 2.2, 3.1, and for $T_r(C, i)$, i being one of 2.1, 2.3, is described in Lemma 3.3 (the same proof holds because Facts 1, 3, 4, and 5 and Lemma 4.7 continue to hold for doubling strips).

We shall consider the remaining case for $T_r(C, i)$ here, involving nondoubling cliques (i.e., i being 2.2). A similar proof will hold for $T_l(C, i)$ with i being one of 2.3, 3.2. We cannot appeal to the proof of Lemma 3.3 directly because Fact 3 does not hold for nondoubling strips.

However, it is easy to see that the fact stated in the beginning of the proof of Lemma 3.3 still holds. We will give the rest of the proof, assuming this fact is true.

Suppose the strips in $T_r(C, i)$ are not right nested. By the above fact, if each strip in $T_r(C, i)$ r -spans the smallest strip in $T_r(C, i)$, then there are no disjoint strips, and the strips in $T_r(C, i)$ must be right nested. So there must exist a strip in $T_r(C, i)$ which is disjoint from the smallest strip A in $T_r(C, i)$; consider the smallest such strip B . Clearly, neither A nor B can r -span any strip in $T_r(C, i)$. We will show that all

other strips D in $T_r(C, i)$ must r-span the lower of A, B . It would then follow from the above fact that the strips in $T_r(C, i)$ with the upper of A, B removed are right nested. We consider the case when B is below A ; the other case is identical.

Consider a strip D as above, and suppose it does not r-span B . By the fact above, it must be disjoint from B ; in addition, it either r-spans A or is disjoint from A as well. Let X, Y denote the upper two strips among A, B, D . So either X and Y are disjoint (as in Figure 10(b)) or (without loss of generality) $l(X) < l(Y)$ and Y r-spans X (as in Figure 10(a)). Let X', Y' be the right successors of X, Y , respectively. Since the right witnesses of A, B, D form a clique, rectangle R containing these witnesses has its lower edge below the lower holes of X, Y and upper edge above the upper holes of X, Y . This, coupled with the fact that X' must stab vertically through R , implies that $l(X') > l(Y), l(X)$; further, if X, Y are indeed disjoint, then $l(X') > l(X) + l(Y) > 2l(X)$. But the latter cannot happen as X is right nondoubling, and therefore X' and X are in the same category, which implies that $l(X') < 2l(X)$. Thus, it must be the case that Y and X are not disjoint; i.e., Y r-spans X . Then $l(X) < l(Y) \leq l(X')$, and Y must be in the same category as X . Since Y r-spans X , is in the same category as X , and is to the left of X' , X' cannot be the right successor of X , a contradiction. \square

Remark. As we mentioned earlier in section 3.2, we can now ignore exception strips from all cliques of OPT .

Next, we generalize Lemma 3.4.

LEMMA 4.9. *Strips in $T_l(C, i)$, i being one of 2.1, 2.2, 2.3, 3.1, 3.2, and in $T_r(C, i)$, i being one of 2.1, 2.2, 2.3, are in distinct size categories.*

Further, consider a right family (left family, respectively) or a set $T_r(C, i)$ of strips, i being one of 2.1, 2.2, 2.3 (set $T_l(C, i)$ of strips, respectively, i being one of 2.1, 2.2, 2.3, 3.1, 3.2) for some clique C in OPT . For any x , the number of strips whose length is more than 2^x times the length of one of the two previous strips to the left (right, respectively) is $O(\frac{\log n}{x})$.

Proof. For the first part, we cannot use the proof of the first part of Lemma 3.4, because that proof was based on Assumption 1, and so we describe it below.

Suppose two strips A, B in $T_l(C, i)$, i being one of 2.1, 2.2, 2.3, 3.1, 3.2, are in the same size category. Without loss of generality, assume B is to the left of A . By Lemma 4.8, B l-spans A . By the definition of a left nondoubling strip, A must be left nondoubling, and its left successor must either lie to the right of B or be B itself. In either case, A 's left witness cell will be independent from B 's left witness cell, a contradiction.

A similar proof holds for two strips A, B in $T_r(C, i)$, i being one of 2.1, 2.2, 2.3.

The second part of the lemma follows by using the same argument as in Lemma 3.4, since this uses only Fact 1 for families, and the first part for cliques. \square

The following lemma generalizes Lemma 3.5.

LEMMA 4.10. *Let $A, B \in T_r(C, i)$ ($T_l(C, i)$, respectively) for some clique C in OPT and some class i , i being one of 2.1, 2.2, 2.3 (2.1, 2.2, 2.3, 3.1, 3.2, respectively), with $l(A) < l(B)$. Let A', B' be the right successors (left successors, respectively) of A, B , respectively. These four strips must be in the following order from left to right (right to left, respectively): A, B, B', A' . In addition, A' cannot r-span (l-span, respectively) B and must have its upper hole above that of B .*

Proof. The proof for $T_l(C, i)$, i being one of 2.1, 2.2, 3.1, and for $T_r(C, i)$, i being one of 2.1, 2.3, is described in Lemma 3.5 (the same proof holds because Facts 1, 3, 4, and 5 continue to hold for doubling strips).

We need to prove the lemma for the remaining cases only. We prove the lemma for $T_r(C, 2.2)$; the other cases have analogous proofs.

From Lemma 4.8, it follows that B r -spans A and hence must lie to the right of A . If B is to the right of A' , then the right witness cells a and b of A and B , respectively, are independent (see Figure 11(b)). Therefore, B is to the left of A' and to the right of A . To show that the right successor of B also lies between A and A' , we show in the next paragraph that A' cannot r -span B . Since A' and B both r -span A , they cannot be disjoint either. Then it follows that the right blocking hole of B is vertically aligned with or to the left of A' . Therefore, the right successor of B is also to the left of A' .

Suppose A' r -spans B . Since we are considering category 2.2, A' and A are in the same category. Since B r -spans A and A' r -spans B , A, B are also in the same category, and then B , and not A' , will be the right successor of A .

It remains to show that the upper hole of A' is above that of B . Recall that B 's right witness cell lies on the horizontal line containing the upper hole of B . If the upper hole of A' is horizontally aligned with or below that of B , the witness cells a and b of A and B , respectively, would be independent (see Figure 11(d)). Hence, the upper hole of A' lies above that of B . \square

DEFINITIONS. As before, we define strip A to be a *right jumper* if its right successor has length at least $2^\Delta l(A)$. Since Δ will be set to at least 1, all right jumpers are actually right doubling strips. Left jumpers are defined analogously.

The following lemma generalizes Lemma 3.6.

LEMMA 4.11. *Consider strips $A, B \in T_r(C, i)$ ($T_l(C, i)$, respectively) for some clique C in OPT and i being one of 2.1, 2.2, 2.3 (2.1, 2.2, 2.3, 3.1, 3.2, respectively). Suppose $l(A) < l(B)$. The following two facts hold.*

1. *The vertical separation between the upper holes of A and B is at most $2^\Delta l(A)$.*
2. *If A is not amongst the smallest $\gamma\Delta$ nonjumper strips in $T_r(C, i)$ ($T_l(C, i)$, respectively), the vertical separation between the upper holes of A and B is at most $\frac{l(B)}{2^{(\gamma-1)\Delta}}$.*

Proof. The proof is identical to that of Lemma 3.6 (the appropriate generalizations of the lemmas used there have to be invoked) because that proof uses only Fact 4 and Lemmas 4.8, 4.9, and 4.10, which hold for both doubling and nondoubling strips. \square

We generalize Lemma 3.7 next.

LEMMA 4.12. *Let $A, B \in T_r(C, i)$ ($T_l(C, i)$, respectively) for some clique C in OPT and i being one of 2.1, 2.2, 2.3 (one of 2.1, 2.2, 2.3, 3.1, 3.2, respectively). Further, suppose $l(A) < l(B)$. Then A lies completely above the right (left, respectively) blocking hole of B .*

Proof. The proof is the same as that of Lemma 3.7 (again, the appropriate generalizations of the lemmas used there have to be invoked), as that proof uses only Facts 4 and 5 and Lemmas 4.8 and 4.10, which continue to hold for both doubling and nondoubling strips. \square

Finally, we need the following lemma, which generalizes Lemma 3.8.

LEMMA 4.13. *Let C and C' be the cliques in OPT containing the right and left witness cells, respectively, of a strip A . Let B be a strip in $T_r(C, i)$ smaller than A . Let B' be strip in $T_l(C', i)$ smaller than A . If i is one of 2.1 or 2.2, then B' cannot l -span B . If i is one of 2.1 or 2.3, then B cannot r -span B' .*

Proof. The proof for $i = 2.1, 2.2$ is identical to that of Lemma 3.8 as Facts 1, 3, 4, and 5 used in that proof continue to hold in the left direction for $i = 2.1, 2.2$ (these

involve left doubling strips). Note that the whole of the proof for Lemma 3.8 is not being invoked; rather, only the first half is being invoked. To invoke the second half (namely, B cannot r-span B') as well, we will need the above facts to hold in the right direction. This is indeed true for Class 2.1, which is right doubling as well but not for Class 2.2. Invoking this second half for Class 2.1, we get the last part of the lemma for this class. The proof for $i = 2.3$ is analogous. \square

4.5. Accounting for strips. The strips in each class are accounted for separately. The number of strips in Class 1.1 is $O(|OPT|)$ by Lemma 4.6. The number of strips in Class 1.2 is $O(|OPT|^{\frac{\log n}{\Delta}})$ by Lemmas 4.6 and 4.9.

Class 2.1 contains strips that are doubling in both directions, like the strips in the special case considered in previous sections. The number of strips in Class 2.1 can be shown to be $O(|OPT|(\Delta + \frac{\log n}{\Delta}))$ by a proof identical to those of Lemma 3.9 and Corollary 3.10; however, invocations of lemmas in section 3.2 need to be modified to point to their respective counterparts in section 4.4 (see the remark in section 3.3).

The reason why the proof of Lemma 3.9 does not extend to other classes as well is that Lemma 3.9 uses Lemma 3.8 in both directions, i.e., to claim that B' cannot r-span B and B cannot l-span B' ; this can be done only for Class 2.1. Therefore, the remaining classes need separate proofs, which are given below.

4.5.1. Classes 2.2 and 2.3.

LEMMA 4.14. *Consider a strip $P \in T_r(C, 2.2)$ and its right follower $Q \in T_r(C, 2.2)$ (Q need not be defined). Then one of the following must hold.*

1. P is either the largest strip or among the smallest $\Delta + 2$ strips in $T_r(C, 2.2)$.
2. Q is among the smallest $4\Delta + 1$ strips in $T_l(C', 2.2)$, where C' denotes the clique in OPT containing the left witness cell of Q .
3. Let P' be the strip in $T_r(C, 2.2)$ which is to the left of P and shorter than P such that the number of strips longer than P' and shorter than P in $T_r(C, 2.2)$ is 3. Then either P' does not exist or $2^{3\Delta}l(P') \leq l(Q)$.
4. Let Q' denote the strip in $T_l(C', 2.2)$ whose left follower is Q . Then $2^\Delta l(Q') \leq l(Q)$.

Proof. We suppose that none of the four conditions hold and derive a contradiction.

Since condition 1 does not hold, Q exists. The proof then proceeds using the following claims, which are proved in subsequent paragraphs. We claim that any strip in $T_r(C, 2.2)$ with length less than $l(Q')/2$ must lie completely above Q' . Since condition 3 is violated, P' must exist. We then show that P' has length less than $l(Q')/2$ and therefore lies above Q' . Since Q r-spans P' by Lemma 4.8, the vertical separation between the upper holes of Q and Q' is at least $l(P')$. By the violation of condition 2 and Lemma 4.11 applied to Q' and Q , the vertical separation between the upper holes of Q and Q' is at most $\frac{l(Q)}{2^{3\Delta}}$. It follows that $l(P') \leq \frac{l(Q)}{2^{3\Delta}}$. This satisfies condition 3, a contradiction.

First, we show that any strip R in $T_r(C, 2.2)$ with length less than $l(Q')/2$ must lie completely above Q' . Clearly, R lies to the left of Q . Its right successor S is to the right of Q , by Lemma 4.10. Further, S has size less than $l(Q')$ (since $l(R) < l(Q')/2$ and Class 2.2 is right nondoubling). In addition, by Lemma 4.10, the upper hole of S is above that of Q . It follows that the lower hole of S must also be above that of Q' ; otherwise, $l(S) \geq l(Q')$, a contradiction. Suppose that R is not completely above Q' . We will get a contradiction as follows. Since R is not completely above Q' and S r-spans R , S is not completely above Q' either. Then one of the two situations shown




FIG. 17. (a) S is to the left of Q' , and Q does not l -span Q' . (b) S is to the right of Q' , and S does not r -span R since it is not above Q' .

in Figure 17 must hold, depending upon whether S is to the left or right of Q' . In the first case, Q cannot l -span Q' , and in the second case S cannot r -span R , both contradictions. Therefore, R is completely above Q' .

Second, we show that P' has length less than $l(Q')/2$. To do this, we will show that $l(P) \leq 2l(Q')$. Then, since there are three strips between P and P' in $T_r(C, 2.2)$, $l(P') < l(P)/4 \leq l(Q')/2$. That $l(P) \leq 2l(Q')$ is shown as follows. We show in the next paragraph that P must have its upper hole aligned with or below that of Q and its lower hole aligned with or above that of Q' . Thus $l(P)$ is at most the vertical separation between the upper hole of Q and the lower hole of Q' . Since condition 2 is violated, Lemma 4.11 applied to Q' and Q implies that the vertical separation between their upper holes is at most $\frac{l(Q)}{23\Delta}$. Thus the vertical distance between the upper hole of Q and the lower hole of Q' is at most $l(Q') + \frac{l(Q)}{23\Delta} < l(Q')(1 + \frac{1}{22\Delta})$, by the violation of condition 4. Thus $l(P) \leq l(Q')(1 + \frac{1}{22\Delta}) \leq 2l(Q')$, as required.

It remains to show that P must have its upper hole aligned with or below that of Q and its lower hole aligned with or above that of Q' . By Lemma 4.8, Q r -spans P , and therefore the upper hole of P is aligned with or below that of Q . By the violation of condition 1 and by Lemma 4.8, which states that all strips in $T_r(C, 2.2)$ are in distinct categories, there exists a strip R in $T_r(C, 2.2)$ such that $l(R) < \frac{l(P)}{2\Delta+1} < \frac{l(Q)}{2\Delta+1}$. Then, by the violation of condition 4, $l(R) < l(Q')/2$. From the earlier part of this proof, it follows that R lies completely above Q' . Clearly, R is to the left of P . By Lemma 4.8, P r -spans R , and therefore the upper hole of P is above that of Q' . Further, by Lemma 4.13, P cannot l -span Q' . Therefore, the lower hole of P must be aligned with or above that of Q' . \square

COROLLARY 4.15. *The number of strips in Class 2.2 is $O(|OPT| * (\frac{\log n}{\Delta} + \Delta))$.*

Proof. We consider four subclasses, depending upon which of the conditions in Lemma 4.14 is satisfied. The number of strips P which satisfy the first condition is clearly $O(|OPT| * \Delta)$ because each clique in OPT has $O(\Delta)$ such strips. The number of strips P which satisfy the third condition is $O(|OPT| * \frac{\log n}{\Delta})$, using arguments similar to those used for conditions 4 and 5 of Lemma 3.9 in the proof of Corollary 3.10. Next, consider strips P such that either condition 2 or 4 holds. Such a strip P has a unique right follower Q in $T_r(C, 2.2)$. Note that any strip in $T_r(C, 2.2)$ is the right follower of at most one strip. Thus it suffices to bound the number of strips Q which are right followers of strips P satisfying condition 2 or 4. Using arguments similar to those in Corollary 3.10, the number of such strips Q satisfying condition 2 can be shown to be $O(|OPT| * \Delta)$, and the number of such strips Q satisfying




FIG. 18. The two scenarios in Lemma 4.17: (a) $j = 2.3, 3.2$. (b) $j = 2.1, 2.2, 3.1$.

condition 4 can be shown to be $O(|OPT| * \frac{\log n}{\Delta})$. \square

A similar argument as above works for Class 2.3.

COROLLARY 4.16. The number of strips in Class 2.3 is $O(|OPT| * (\frac{\log n}{\Delta} + \Delta))$.

4.5.2. Classes 3.1 and 3.2.

LEMMA 4.17. Consider a strip P in $T_l(C, 3.1)$ or $T_l(C, 3.2)$. Let Q be the right successor of P . Let j be the class containing strip Q and C' be the clique containing the left witness cell of Q . Then either $j = 1.1$ or $j = 1.2$, or Q is among the smallest $2\Delta + 2$ strips in $T_l(C', j)$.

Proof. Since P belongs to Class 3.1 or Class 3.2 and not to Class 1.1, it has a well-defined right successor Q . Further, since P does not define a right witness cell, the upper holes of P and Q are horizontally aligned (see Figure 18). In addition, if P is right nondoubling, then $l(Q) < 2l(P)$, and if P is right doubling, then $l(Q) \leq l(P)2^\Delta$, as P is not a right jumper (i.e., it is not in Class 1.2). Note that if P is in Class 3.2, then it must be right doubling, as all strips in \mathcal{S} are either left doubling or right doubling or both (recall the definition of \mathcal{S} from section 4). We consider various cases depending upon the nature of Q .

Suppose $j \neq 1.1$ and $j \neq 1.2$. Then Q is not a left terminal strip. Let $T_l(C', j)$ have k strips smaller than Q . Let these be R_1, \dots, R_k , in increasing order of length. We need to show that $k \leq 2\Delta + 1$.

By Lemmas 4.8 and 4.9, the R_i s and Q together form a left nested set of strips and therefore belong to distinct categories. Note that since $j \neq 1.1$, Q and each of the R_i s have left successors. Let R'_i denote the left successor of R_i . There are two cases now, depending upon whether j is one of 2.3, 3.2 or one of 2.1, 2.2, 3.1.

First, suppose j is one of 2.3, 3.2. Then Q and the R_i s are all left nondoubling (see Figure 18(a)). Then R'_i is in the same category as R_i . Since all R_i s and Q are left nested and in distinct size categories, $l(R'_1) < l(R'_2) < \dots < l(R'_k) < l(Q)$, and $2^{k-1}l(R'_1) < l(Q)$. All R'_i s must be between P and Q . For R'_i cannot be to the right of Q by Lemma 4.10. And, if R'_i is to the left of P , then the left witness cells of Q and R_i are independent because Q 's left witness cell is on the horizontal line joining the upper holes of P and Q . From Lemma 4.10, the upper hole of each R'_i is above the upper holes of both Q and P . Each R'_i must r-span P because R'_i must l-span R_i and the hatched region is hole-free (because Q r-spans P). Thus $l(P) < l(R'_1)$. Therefore $l(Q) > 2^{k-1}l(R'_1) > 2^{k-1}l(P)$. So if $k \geq \Delta + 1$, $l(Q) > l(P)2^\Delta$, a contradiction (see the first paragraph of this proof; note that $\Delta \geq 1$). It follows that $k \leq \Delta$ in this case.

Second, suppose j is one of 2.1, 2.2, 3.1. Q and the R_i s are all left doubling (see Figure 18(b)). Note that the hatched regions in the figure must be hole-free, as Q must r -span P and l -span R_i . Then the lower hole of each R_i must be below that of P ; otherwise, the left witness cell for R_i will be on or to the left of P and independent from the left witness cell for Q , a contradiction. Therefore, the vertical distance between the upper hole of P (or of Q) and the lower hole of R_1 is at least $l(P)$. Since the R_i s and Q belong to distinct categories, $2^\Delta l(P) \geq l(Q) > 2^{k-1}l(R_1)$ (the first inequality follows from the first paragraph of this proof). It follows that the vertical separation t between the upper holes of R_1 and Q is at least $l(P) - l(R_1) > (2^{k-1-\Delta} - 1)l(R_1)$. For $k \geq 2\Delta + 2$, $t > (2 \cdot 2^\Delta - 1)l(R_1) > 2^\Delta l(R_1)$. This contradicts Lemma 4.11 (the first part, applied to R_1 and Q). Thus $k \leq 2\Delta + 1$ in this case, as required. \square

COROLLARY 4.18. *The number of strips in Classes 3.1 and 3.2 is $O(|OPT| * (\frac{\log n}{\Delta} + \Delta))$.*

Proof. Each strip in these two classes has a right successor, which in turn has a left witness by Lemma 4.4. Further, by Lemma 4.5, any strip is the right successor of at most two strips. By Lemma 4.17, either (a) the right successor of a strip in these two classes is in Class 1.1 or 1.2, or (b) the right successor of a strip in these two classes is in some class $j \neq 1.1, 1.2$ and is amongst the smallest $2\Delta + 2$ strips in $T_i(C', j)$ for some clique C' in OPT . Strips in Classes 3.1 and 3.2 for which the right successor satisfies the latter property are clearly $O(|OPT| * \Delta)$ in number. And strips for which the right successor satisfies the former property are $O(|OPT| \frac{\log n}{\Delta})$ in number by Lemmas 4.6 and 4.9. \square

4.5.3. Summing up.

THEOREM 4.19. *The number of rectangles needed to cover the given polygon is $\Omega(\#N/\sqrt{\log n})$, where $\#N$ is the number of necessary strips and therefore the number of rectangles used by our algorithm.*

Proof. From Corollaries 3.10, 4.15, 4.16, and 4.18, it follows that $|\mathcal{S}| = O(|OPT| * \max\{\frac{\log n}{\Delta}, \Delta\})$. By Lemma 4.3, $|\mathcal{S}|$ is at least a quarter of the number of rectangles used by our algorithm, which is equal to the number of necessary strips. The theorem follows by setting $\Delta = \sqrt{\log n}$. \square

5. Counterexample. If the average family size was $O(1)$ for all polygons, then our claim that there always exist $\#F$ independent points will give a constant factor approximation algorithm. Unfortunately, this is not the case. Here we give an example of a polygon in which the average family size is $\theta(\frac{\log n}{\log \log n})$.

This example can also be slightly modified in a way such that all right witness points can be covered by $O(\frac{\#N \log \log n}{\log n})$ cliques and the average right clique size is $\Theta(\frac{\log n}{\log \log n})$, but the average family size (both left and right) remains $\Theta(\frac{\log n}{\log \log n})$. However, in this example, covering the left witness points requires $\Omega(\#N)$ cliques, and the maximum clique size is $O(1)$ for these points. This example was the key to our lower bound. We do not know whether there are examples where the average clique size is superconstant for both the left witness points and the right witness points. In this sense, our bound of $\sqrt{\log n}$ does not seem like an unnatural meeting point.

Let l be a parameter, which we will ultimately set to $\log n$. We give an example where the average left family size, average right family size, and average right clique size are all $\Theta(\log_l n)$.

Our polygon will have two kinds of holes, *nonblocking* and *blocking*. There will be $\Theta(n)$ blocking holes and $\Theta(n \log_l n)$ nonblocking holes; so most holes will be non-




FIG. 19. Arrangement of columns of sparsity i and $i + 1$.

blocking. First, we will describe the arrangement of nonblocking holes and then that of blocking holes. All subsequent references to strips will be to those formed by nonblocking holes.

Nonblocking holes are arranged in columns. Each column has a certain *sparsity*. A column with sparsity i will have $\frac{n}{l^i} + 1$ holes in it, where $0 \leq i \leq \log_l n$; these holes will be put in rows $0, l^i, 2l^i, \dots$. So the least sparse column will have $n + 1$ holes in rows $0, 1, 2, 3, \dots, n$, and the most sparse will have 2 holes in rows $0, n$. There will be $l^{i-1}(l + 1)$ columns of sparsity i , $i \geq 1$. Therefore, the total number of nonblocking holes will be $\Theta(n \log_l n)$. The arrangement of these columns can be described by the following sequential procedure.

The leftmost and rightmost columns will have sparsity 0. Between these two columns, put $l + 1$ columns of sparsity 1; these $l + 1$ columns together constitute a *pack*. Then, between each pair of consecutive columns of sparsity 1, put a pack of $l + 1$ columns of sparsity 2, and so on, as shown in Figure 19. Note here that a pack of sparsity $i + 1$ columns is put only between pairs of consecutive columns of sparsity i which belong to the same pack; these sparsity i columns will not have any columns of sparsity less than i between them.

Blocking holes will always be placed as follows. First, we form right families comprising strips formed by nonblocking holes in columns which are not the last in their respective packs. Note that most (all but $\Theta(n)$) nonblocking holes lie in such columns. Consider a column C with sparsity i which is not the rightmost column in its pack. There are $\frac{n}{l^i}$ strips in such a column. These strips are organized into *groups* of l strips each, the strips in each group being vertically consecutive. Consider one such strip which is the j th strip in its group. We define a right successor s' for s , where s' is the unique strip in the column C' defined below which r -spans s ; C' is the j th leftmost column amongst the pack of $l + 1$ sparsity $i + 1$ columns nested between C and the next sparsity i column to the right of C . The size of each right family defined by the above right successors is clearly large, i.e., $\Theta(\log_l n)$.

We will now arrange blocking holes so that all strips in each right family defined above will indeed have a common right blocking hole. For each right family defined above, put a blocking hole immediately to the right of the rightmost strip in such a way that it blocks all strips in this family. The number of blocking holes put is clearly $\Theta(n)$.

Thus, what we have achieved above is an arrangement of $\Theta(n \log_l n)$ holes where all but $O(n)$ strips lie in right families of size $\Theta(\log_l n)$. We remark here that all but $O(n)$ strips lie in left families of size $\Theta(\log_l n)$ as well in this arrangement. The picture




FIG. 20. A basic block.

of a “basic block” shown in Figure 20 will be helpful in seeing that this is true.

A basic block comprises the following holes.

1. holes bounding a group of strips on a column C of sparsity i , where C is not the rightmost column in its pack. Let h, h' be the topmost and bottommost such holes (see Figure 20);
2. all blocking holes whose vertical position is between h and h' and whose horizontal position is between C and the next column in the pack containing C ;
3. all nonblocking holes which are located vertically between h and h' and are on columns of sparsity $i + 1$ between C and the next column in the pack containing C .

5.1. Getting large right cliques. We need to make a modification to the above construction to get large right clique sizes while leaving left and right family sizes as before.

The modification is that columns containing nonblocking holes need to be shifted downwards by varying amounts while maintaining most (but not all) of the right and left families as such. This shifting is described by the following sequential procedure.

The shifting procedure is carried out in rounds. In the i th round, only columns of sparsity i or more will be shifted. Assume that the procedure has already been

executed for i rounds. At this point, for any column C of sparsity i , all columns of sparsity $i + 1$ or more which appear between C and the next column in its pack will not have experienced any shift relative to C . For each strip s in a column C of sparsity i which is not the rightmost in its pack, let $map(s)$ denote that strip which has the same right blocking hole as s and lies on a column of sparsity $i + 1$ between C and the next sparsity i column in C 's pack. It can be verified from the shifting procedure below that $map(s)$ is always well defined and that a basic block (such as the one shown in Figure 21) has blocking holes for strips in C forming a group distributed diagonally (this basic block requires forming groups of l consecutive strips on column C , leaving the first $l - 1$ strips out of this grouping; this is to account for the shifts made to C so far). So each basic block formed by groups on C looks like the one in Figure 21, except that holes in this basic block on sparsity $i + 1$ columns between C and the next sparsity i column D in C 's pack are all aligned with h , the top hole of this group. The $i + 1$ st round proceeds as follows.

For each column C of sparsity i which is not the rightmost column in its pack, consider any group of strips on C . Let h be the topmost hole in this group. Each strip s in this group is considered in turn. Let \mathcal{P} denote the pack of sparsity $i + 1$ strips nested between C and the next sparsity i strip to the right. The column C'' immediately preceding the column C' containing $map(s)$ in \mathcal{P} is shifted down so that the hole which was horizontally aligned with h is now aligned with the upper hole of s (see Figure 21). In addition, all columns nested between C'' and C' will also be shifted down so that no relative shift is introduced between C'' and these columns in this round. For future reference, we denote the strip on C'' which now r-spans s by $cmap(s)$ and the strip on C' which now r-spans s by $newmap(s)$. Note that $cmap(s)$ is defined unless s is the first strip in its group. Figure 21 shows a basic block after this modification.

The above shifting procedure modifies right families, because right successors of strips could have changed. For each strip s in C , the right successor changes from $map(s)$ to $newmap(s)$. Families defined by this new definition of right successor are also large, essentially because a right successor can be defined for every strip other than those which are on the last columns in their packs. Thus, all but $O(n)$ of the strips will continue to be in right families of size $\Theta(\log_l n)$.

Also, the average right clique size is $\Theta(\log_l n)$. To see this, note that the right witness points of $s, cmap(s), cmap(cmap(s)), \dots$ form a right clique and that $cmap(s)$ is defined for all those s which are not the first strips in their respective groups or in the last columns in their respective packs. Since there are only $O(n)$ strips s which are either the first strips in their respective groups or in the last columns in their respective packs, the total number of right cliques is $\Theta(n)$ and the average right clique size is $\Theta(\log_l n)$.

It now remains to show that left families continue to be large after the above modification. Consider a column C of sparsity i and the next column D to its right in its pack. As is clear from Figure 21, left successors can be defined in the pack \mathcal{P} for each of the strips in D , except those strips which are either the first or last in their respective groups. Defining left successors recursively in this way ensures that the average left family size is also $\Theta(\log_l n)$.

5.2. Large left and right cliques? In the above example, it can be seen with some effort that all left cliques have size $O(1)$. We do not know whether this example can be modified so that the average left and right cliques sizes are both large.




FIG. 21. Two basic blocks obtained after the $i + 1$ st round of shifting. White holes are blocking holes. One large clique has been highlighted. The crosses are right witness points.

6. Conclusions. A number of loose ends remain for this problem. The main question, of course, is whether the approximation factor can be brought down to $O(1)$. Another question is whether there exists a polygon whose clique cover and independent set numbers are small-o of the number of necessary strips.

Related problems. We briefly mention some related problems and the current state of knowledge on these problems.

Non-axis-parallel rectangles. One variant of the above rectangle covering problem is when the covering rectangles need not be axis-parallel.

Our techniques do not seem to extend to this case. However, they do extend




FIG. 22. Covering with non-axis-parallel rectangles.

even when all the covering rectangles must be inclined at the same angle or at one of a constant number of angles. But there are examples where rectangles inclined at an arbitrary number of angles are involved in the optimal cover. And, in this case, the issue seems to be different and related to the problem of covering a given set of points using a minimum number of straight lines. We do not know the exact nature of this relationship though. No $o(\log n)$ approximation factor is known for this problem either (see [1]). We describe an example below where the optimum cover has size $O(n)$ when non-axis-parallel rectangles are allowed, whereas it is $\Omega(n\sqrt{n})$ if only axis-parallel rectangles are allowed. This gives support for the intuition that the size of the minimum cover should be much smaller if non-axis-parallel rectangles are allowed.

Partition the $n \times n$ grid into a^2 tiles of size n^2/a^2 each (as in Figure 22(a)). One such tile is shown in Figure 22(b). Each tile has dimensions $n/a \times n/a$. The structure of the tiles results in a partition of the grid into triangles along the boundary and rhombuses inside. Each rhombus has n/a holes on each of its sides. First, consider the case when only axis-parallel rectangles are allowed. Each rhombus needs n/a axis-parallel rectangles for covering. Similarly, each triangle needs n/a rectangles to be covered. The total number of triangles is $4a$, and the total number of rhombuses is $\Theta(a^2)$. Therefore, the optimum has size $\Theta(a^2n/a + 4an/a)$. Next, consider the case when arbitrarily oriented rectangles are allowed. Now each rhombus can be covered by just one rectangle. Therefore the cover size is $\Theta(an/a + 4an/a)$. For $a = \sqrt{n}$, the cover sizes are $\Theta(n\sqrt{n})$ and $\Theta(n)$, respectively.

Nonrectilinear polygons. When the polygon itself is not rectilinear but has only obtuse angles, suitably discretizing the problem so as to apply the greedy set covering algorithm [10] is itself nontrivial. Levcopoulos and Gudmundsson [14] showed that this can indeed be done. So this problem too has an $O(\log n)$ factor approximation algorithm, and no better bound is known.

Rectilinear polygons and fat rectangles. When the covering objects are squares or rectangles with bounded aspect ratio, then Levcopoulos and Gudmundsson [15] give a constant factor approximation algorithm.

Acknowledgment. We thank the referees for extensive comments that have helped to improve the presentation.

REFERENCES

- [1] V.S. ANIL KUMAR, S. ARYA, AND R. HARIHARAN, *Hardness of set covering with intersection 1*, in Proceedings of the 27th International Colloquium on Automata, Languages and Programming, Geneva, Switzerland, 2000, pp. 624–635.
- [2] P. BERMAN AND B. DASGUPTA, *Approximating rectilinear polygon cover problems*, *Algorithmica*, 17 (1997), pp. 331–356.
- [3] H. BRÖNNIMANN AND M. GOODRICH, *Almost optimal set covers in finite VC-dimension*, *Discrete Comput. Geom.*, 14 (1995), pp. 263–279.
- [4] S. CHAIKEN, D.J. KLEITMAN, M. SAKS, AND J. SHEARER, *Covering regions by rectangles*, *SIAM J. Alg. Disc. Meth.*, 2 (1981), pp. 394–410.
- [5] Y. CHENG, S.S. IYENGAR, AND R.L. KASHYAP, *A new method for image compression using irreducible covers of maximal rectangles*, *IEEE Trans. Software Engrg.*, 14 (1988), pp. 651–658.
- [6] J.C. CULBERSON AND R.A. RECKHOW, *Covering polygons is hard*, *J. Algorithms*, 17 (1994), pp. 2–44.
- [7] D.S. FRANZBLAU, *Performance guarantees on a sweep-line heuristic for covering rectilinear polygons with rectangles*, *SIAM J. Disc. Math.*, 2 (1989), pp. 307–321.
- [8] D.S. FRANZBLAU AND D.J. KLEITMAN, *An algorithm for constructing regions with rectangles*, *Inform. and Control*, 63 (1984), pp. 164–189.
- [9] A. HEGEDÜS, *Algorithms for covering polygons with rectangles*, *Comput. Aided Geom. Design*, 14 (1982), pp. 257–260.
- [10] D.S. JOHNSON, *Approximation algorithms for combinatorial problems*, *J. Comput. System Sci.*, 9 (1974), pp. 256–278.
- [11] J.M. KEIL, *Minimally covering a horizontally convex orthogonal polygon*, in Proceedings of the 2nd Annual ACM Symposium on Computational Geometry, Yorktown Heights, NY, 1986, pp. 43–51.
- [12] C. LEVCOPOULOS, *A fast heuristic for covering polygons by rectangles*, in Proceedings of Fundamentals of Computer Theory, Lecture Notes in Comput. Sci. 199, Springer, Berlin, 1985, pp. 269–278.
- [13] C. LEVCOPOULOS, *Improved bounds for covering general polygons by rectangles*, in Proceedings of 6th Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Comput. Sci. 287, Springer, Berlin, 1987, pp. 95–102.
- [14] C. LEVCOPOULOS AND J. GUDMUNDSSON, *Close approximations of minimum rectangular covering*, in Proceedings of 16th Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Comput. Sci. 1180, Springer, Berlin, 1996, pp. 135–146.
- [15] C. LEVCOPOULOS AND J. GUDMUNDSSON, *Approximation algorithms for covering polygons with squares and similar problems*, in Proceedings of RANDOM'97, Lecture Notes in Comput. Sci. 1269, Springer, Berlin, 1997, pp. 27–41.
- [16] L. LOVASZ, *On the ratio of optimal integral and fractional covers*, *Discrete Math.*, 13 (1975), pp. 383–390.
- [17] R. RAZ AND S. SAFRA, *A sub-constant error-probability low-degree test and a sub-constant error-probability PCP characterization of NP*, in Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing, El Paso, TX, 1997.
- [18] C. LUND AND M. YANNAKAKIS, *On the hardness of approximating minimization problems*, in Proceedings of the 25th Annual ACM Symposium on Theory of Computing, San Diego, CA, 1993, pp. 286–293.
- [19] W.J. MASEK, *Some NP-Complete Set Covering Problems*, manuscript, MIT, Cambridge, MA, 1979.