Parallel Two Dimensional Witness Computation*

Richard Cole † *Zvi Galil*‡ *Ramesh Hariharan* § *S. Muthukrishnan* ¶
Kunsoo Park ‖

May 18, 2001

Abstract

An optimal parallel CRCW-PRAM algorithm to compute witnesses for all non-period vectors of an $m_1 \times m_2$ pattern is given. The algorithm takes $O(\log\log m)$ time and does $O(m_1 \times m_2)$ work, where $m = \max\{m_1, m_2\}$. This yields a work optimal algorithm for 2D pattern matching which takes $O(\log\log m)$ preprocessing time and $O(1)$ text processing time.

1 Introduction

We consider the problem of computing witnesses in parallel for all non-period vectors of a pattern p with m_1 rows and m_2 columns.

The significance of this problem is that all known optimal alphabet-independent parallel and sequential algorithms for 2D pattern matching [ABF92, ABF93, GP92, CGHMR92] require precomputation of witnesses for all non-period vectors of row length less than $\frac{m_1}{k}$ and column length less than $\frac{m_2}{k}$, where the value of the constant k affects the text processing time and work by only a constant factor. Here, the row length of a vector is the number of rows it spans; column length is defined analogously.

Witness computation can be accomplished easily in $O(m_1 \times m_2 \log|\Sigma|)$ time using suffix trees, where $|\Sigma|$ is the size of the alphabet from which the pattern characters are drawn. The first linear time (i.e., $O(m_1 \times m_2)$) sequential algorithm to compute witnesses for all non-period vectors of p of row length less than $\frac{m_1}{4}$ and column length less than $\frac{m_2}{4}$ was obtained recently by Galil and Park [GP92, GP93]. The only known parallel algorithm for witness computation used suffix trees and was not optimal; it took $O(\log m)$ time using $O(m_1 \times m_2)$ processors on a CRCW-PRAM and $O(\log^2 m)$ time using $O(\frac{m_1 \times m_2}{\log m})$ processors on a CREW-PRAM [AB92], where $m = \max\{m_1, m_2\}$. We give an algorithm which takes $O(\log\log m)$

*A preliminary version of this paper appeared in the 34th IEEE Symposium on Foundations of Computer Science, 1993.

†Courant Institute, NYU. Work supported by NSF grants CCR-8902221, CCR-8906949 and CCR-9202900.

‡Columbia University and Tel-Aviv University. Work partially supported by NSF grant CCR-90-14605 and CISE Institutional Infrastructure Grant CDA-90-24735.

§Courant Institute, NYU. Work supported by NSF grants CCR-8902221, CCR-8906949 and CCR-9202900. Present Address: Indian Institute of Science, Bangalore.

¶Courant Institute, NYU. Work supported by NSF/DARPA grant CCR-89-06949 and NSF grant CCR-91-03953. Present Address: AT&T Labs-Research, Florham Park.

‖Department of Computer Engineering, Seoul National University, Seoul 151-742, Korea.

time and $O(m_1 \times m_2)$ work to compute witnesses for *all* non-period vectors of the pattern on a CRCW-PRAM. When combined with [CGHMR92] and [CGGPR93], our result provides a work optimal algorithm for 2D pattern matching; it takes $O(\log\log m)$ preprocessing time and $O(1)$ text processing time. The lower bound of [BG91] for string matching implies that any algorithm for 2D pattern matching requires $\Omega(\log\log m)$ time using a linear number of processors in the comparison model of computing. In fact, this lower bound also holds for square 2D pattern matching as we show in this paper.

The algorithm has two steps. In the first step, witnesses are computed for all non-period vectors of row length less than $\lfloor \frac{m_1}{16} \rfloor$ and column length less than $\lfloor \frac{m_2}{16} \rfloor$. In the second step, witnesses are computed for all non-period vectors. The second step is an extension of the algorithm for the first step.

A natural approach to performing these steps is to follow the approach for computing string witnesses [BG90, ABG92]. There, witnesses for a string s of length m are found in $O(\log\log m)$ stages. Stage i finds witnesses for a prefix of s of length l_i; these witnesses are used in Stage $i+1$ to find witnesses for the prefix of s of length l_{i+1} in constant time. The lower bound of [BG91] shows that $\Omega(\log\log m)$ stages are required when only $O(m)$ work is allowed.

A natural analog of a prefix of a string is a corner sub-block of the 2d pattern p. The string witness computation algorithm described above is generalized to 2d patterns by designing an algorithm which performs Stage i, i.e., finds witnesses for all non-period vectors of a corner sub-block L_{i+1} of p, given witnesses for all non-period vectors of a smaller corner sub-block L_i of p, for appropriate L_i, L_{i+1}. This presents two problems.

Problem 1. Using a construction similar to the string matching lower bound of [BG91], it is possible to show that, if no additional information is at hand, Stage i takes $\Omega(\log\log |L_i|)$ time even for the simple case in which all period vectors of L_i are parallel. This remains true if L_i is a central rather than a corner sub-block of L_{i+1}, as in the sequential algorithm of [GP92, GP93]. Since there are $\Theta(\log\log m)$ stages, this approach leads to an algorithm which takes $\Omega((\log\log m)^2)$ time even for this simple case. But our goal is an $O(\log\log m)$ time algorithm.

Problem 2. For the general case, i.e., when the period vectors of L_i are non-parallel, it is not clear how to perform Stage i in $O(\log\log m)$ time while preserving work optimality of the overall algorithm. The sequential algorithm of [GP92, GP93] performs Stage i by constructing a series of strings in sequence and then applies a sequential algorithm for finding leftmost witnesses in each string. It is not clear how to parallelize the construction of these strings in $O(\log\log m)$ time, even though leftmost witnesses can be found for each string optimally in $O(\log\log m)$ time and linear work [GaP].

Our $O(\log\log m)$ algorithm has the following two key features.

First, Problem 2 is solved in $O(\log\log m)$ time and $O(|L_{i+1}|)$ work for the general case. The solution is based on a new periodicity property of 2d patterns which we prove and use. Note that this immediately gives an $O((\log\log m)^2)$ time algorithm.

Second, to achieve $O(\log\log m)$ time, we compute with fringes instead of corner or central blocks which circumvents Problem 1. In the main step, witnesses are found for all non-period vectors of p', the central half of p, in $O(\log\log m)$ stages. These stages computes witnesses for all non-period vectors of fringes of p' of successively increasing sizes. Using fringes allows each stage to be performed in $O(1)$ time while keeping the overall work linear. A notable

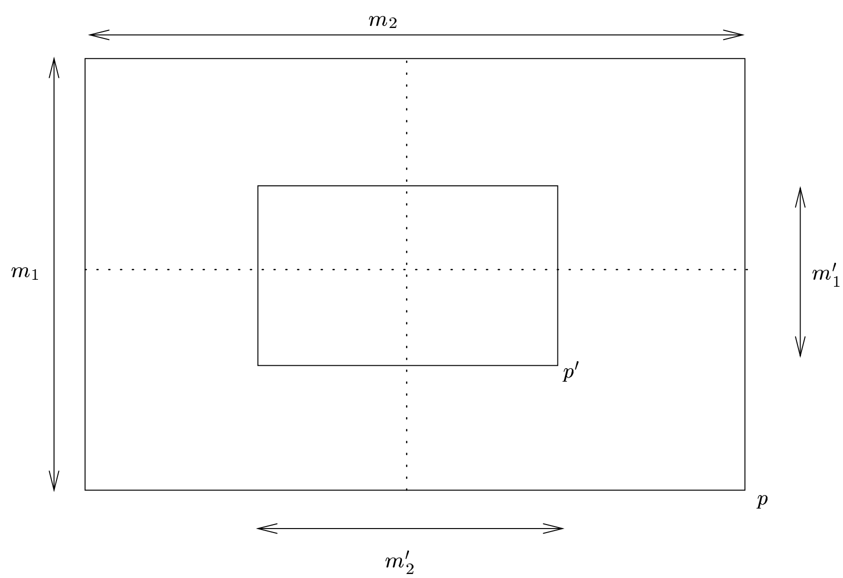

Figure 1: p and p'.

feature of the algorithm is that the witnesses it computes need not be within p'. Finally, using the solution to Problem 2, witnesses in p, if any, are found for the period vectors of p' in $O(\log\log m)$ time.

Section 2 gives some preliminary definitions. Section 3 outlines the main steps of the algorithm for computing witnesses for all non-period vectors of row length less than $\lfloor\frac{m_1}{16}\rfloor$ and column length less than $\lfloor\frac{m_2}{16}\rfloor$. These steps are elaborated upon in Sections 4 and 5. Section 6 gives some useful lemmas required in proving the lemmas in Sections 3, 4 and 5. The proofs of the lemmas in Sections 3, 4 and 5 appear in Sections 7 and 8. Section 9 describes the algorithm for computing witnesses for all non-period vectors of p. Appendix I describes some of the procedures used in the above sections. Appendix II describes the new property of 2d patterns which plays a key role in the algorithm. Though this property appears implicitly in Sections 3, 5 and 9, we believe that it may be of independent interest and therefore deserves an explicit description. Appendix III describes a simple $\Omega(\log\log m)$ time lower bound with a linear number of processors for square 2D pattern matching.

2 Preliminaries

Let p be a size $m_1 \times m_2$ pattern. Let $m = \max\{m_1, m_2\}$. Without loss of generality assume that $m_2 \geq m_1$, i.e., $m = m_2$.

Let m'_1 and m'_2 be the largest even numbers smaller than or equal to $\frac{m_1}{2}$ and $\frac{m_2}{2}$, respectively. Let $m' = \max\{m'_1, m'_2\} = m'_2$. Let p' be the size $m'_1 \times m'_2$ sub-block of p concentric with p (see Fig.1), with ties for the center of p being broken arbitrarily.

Consider two copies of p. Place the second copy so that its left side is aligned with or to the right of the left side of the first copy. In addition, either the top left corner or the bottom left corner of the second copy must be within the first copy. The vector v joining the top left

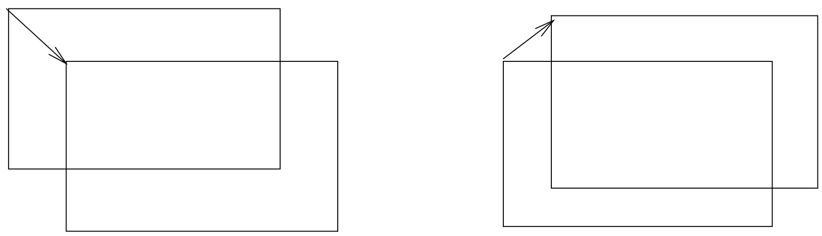

A Quad I Vector A Quad II Vector

Figure 2: Quad I and II vectors.

corner of the first copy to the top left corner of the second copy is called a Quad I vector in the former case and a Quad II vector in the latter case (Fig.2).

Note that a horizontal vector pointing rightwards is both a Quad I and a Quad II vector; in addition, if v is a vertical Quad I vector then $-v$ is a vertical Quad II vector.

The arrow end of a vector is called its *head* and the other end is called its *tail.*

The term "a character a" refers both to a character and its location in the pattern; thus, given a vector v, $a+v$ refers to the location and the character at the location at which the head of v falls when its tail is at the location of character a. We use '$=$' for equality between locations and '$\equiv$' for equality between characters.

The *bottommost* location among a certain set of locations is the bottommost location in this set with ties broken arbitrarily. The *leftmost bottommost* location among a certain set of locations is the bottommost location in this set with ties broken by choosing the leftmost eligible location. *Topmost*, *leftmost*, *rightmost*, *leftmost topmost*, *rightmost bottommost* etc. are defined analogously. In order to determine such locations, we will repeatedly use the algorithm of [FRW88] for finding the leftmost '1' in a binary string s without mention. This algorithm requires constant time and $O(|s|)$ work.

A witness for a vector v is a pair of distinct characters a and $a+v$ (see Fig.3). a is called the *tail* of this witness and $a+v$ is called the *head.* A sub-block of p is said to have a witness for vector v if v has a witness with both head and tail in that sub-block.

Suppose v_1, v_2 are two vectors such that the vector $v_2 - v_1$ has a witness with tail a and head b. If there exists a position c satisfying $c+v_1 = a$ and $c+v_2 = b$ then a witness for either v_1 or v_2 can be found by comparing c with a (see Fig.3). If $c \equiv a$ then $c \not\equiv b$ and a witness for v_2 is found; otherwise, if $c \not\equiv a$ then a witness for v_1 is found. Similarly, if there exists a position c satisfying $c - v_1 = b$ and $c - v_2 = a$, then a witness for either v_1 or v_2 can be found by comparing c with a. This act of finding a witness for one of two vectors using a witness for their difference vector is called a *duel* [Vi85] between the two vectors.

Two overlapping copies of p are said to be *consistent* if the vector joining their top left corners lacks a witness in p; these two copies are then said to be *period overlaps* of each other. The vector joining the top left corners of these two consistent pattern copies is called a *period vector* of p.

The *row length*, $rl(v)$, of a vector v is defined to be the difference of the row coordinates of its end points. *Column length*, $cl(v)$, of v is defined analogously. The *length* of a vector is

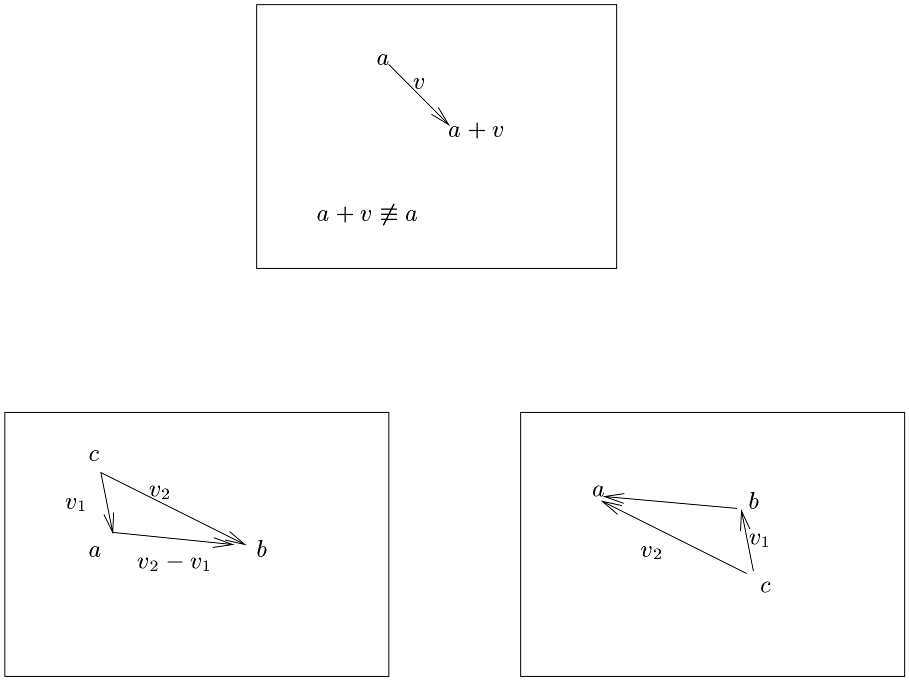

Figure 3: Witnesses and Duelling.

the larger of its row and column lengths. A Quad I or Quad II vector is called *valid* if its row length is less than $\frac{m'_1}{8}$ and its column length is less than $\frac{m'_2}{8}$. Note that $\frac{m'_1}{8} \geq \frac{2\lfloor \frac{m_1}{4} \rfloor}{8} \geq \lfloor \frac{m_1}{16} \rfloor$ and that $\frac{m'_2}{8} \geq \frac{2\lfloor \frac{m_2}{4} \rfloor}{8} \geq \lfloor \frac{m_2}{16} \rfloor$.

Let v_1 and v_2 both be either Quad I or Quad II vectors. v_1 is said to be *clockwise* with respect to v_2 if the head of v_1 appears in the clockwise direction from the head of v_2 when their tails are coincident.

A vector is said to be *eliminated* when a witness for it is found or it is determined that there is no witness for it. A vector is said to *survive* if a witness for it has not yet been found and neither has it been determined whether or not an witness for it exists.

We assume the Common CRCW-PRAM model, i.e., simultaneous writes to the same location by several processors are guaranteed to be of the same value [Ja91].

Primitive $Line(V, R, R')$**.** We need a primitive called $Line(V, R, R')$. Here, R, R' are sub-blocks of p, R is a sub-block of R', and V is some subset of the set of Quad I vectors (Quad II vectors, respectively) with the following property. There exists a line l in R such that the heads of the vectors in V fall on l when their tails are at the top left cornet of R'; in addition, l has the slope of a Quad I vector (Quad II vector, respectively) (see Fig.15). $Line(V, R, R')$ finds witnesses with heads in R and tails in R', if any, for vectors in V in $O(\log \log m)$ time and $O(|R|)$ work. The procedure itself is described in Appendix I.

We assume that $m_1 \geq 16$. If $m_1 < 16$ then witnesses for all Quad I vectors are found by invoking $Line(V_r, p, p)$ for each row index r, $1 \leq r \leq m_1$, where V_r comprises the set of Quad I vectors with row length $r - 1$. This takes $O(\log \log m)$ time and $O(m_1 \times m_2)$ work. Witnesses for all Quad II vectors are found analogously in this case.

We assume that candidate period vectors are stored in a boolean array CP of size $\lfloor \frac{m'_1}{8} \rfloor \times \lfloor \frac{m'_2}{8} \rfloor$; vector v is stored at array location $< rl(v), cl(v) >$. When a process discovers a witness for a vector, it marks the corresponding location in CP accordingly. To process subsets of surviving vectors, all vectors of the appropriate size are considered. Only those which survive are actually processed, but the analysis and the processor allocation proceed as if all survived.

3 The Main Algorithm

We outline the algorithm for computing witnesses for all valid non-period vectors. Once this is done, we will show how to compute witnesses for all non-period vectors in Section 9. The latter uses the same basic ideas as for computing witnesses for valid vectors; however, there are many technicalities due to which it needs to be described separately.

The algorithm for computing witnesses for valid non-period vectors has 4 main steps, Steps A, B, C and D. We first outline these steps. Subsequently, we will describe each step in detail.

Step A. All valid horizontal and vertical vectors are considered and witnesses for them in p', if any, are found.

Step B. All surviving valid Quad I and Quad II vectors are considered in this step. If v and w are two such vectors and $w - v$ is a valid horizontal or vertical vector having a witness in p', then a witness for either v or w is found in this step.

Step C. All surviving valid Quad I and II vectors are considered. For each such vector v, if it has a witness entirely within p' then a witness for v is found. This witness is assured to lie within p but not necessarily within p'. Further, even if v does not have a witness within p', a witness for v may still be found.

Step D. For all surviving valid Quad I and II vectors, witnesses in p, if any, are found.

In the algorithm above, Steps A and B serve to thin the set of survivors; two vectors surviving these steps will have the property that their difference vector, if horizontal or vertical, will have no witnesses in p'. Step C and D complete the job and are the key steps.

We give an outline of each of the four steps next, highlighting the obstacles in generalizing parallel string computation algorithms and sequential 2D witness computation algorithms to parallel 2D pattern matching. Steps A will be described completely. Steps B, C and D will require further elaboration, which will follow in later sections.

3.1 Step A

This step simply involves finding witnesses along each row and column of p'; this is done in $O(\log \log m)$ time and $O(m_1 \times m_2)$ work using the string witness computation algorithm of [BG90, ABG92].

3.2 Step B

This step uses a direct extension of parallel string matching techniques. All surviving valid Quad I and Quad II vectors are considered. Consider the Quad I vectors; Quad II vectors will be processed similarly. Two substeps are performed, each taking $O(\log\log m)$ time and $O(m_1 \times m_2)$ work. Before describing them, we need the following definitions.

Definitions. p' is said to be *r-oriented* if it has a valid horizontal period vector. A vector is said to *fall* at a location c if its head is at c when its tail is coincident with the top left corner of p'. In addition, it is said to *fall* in a sub-block T of p' if $c \in T$.

Substep B.1. All sets of vectors which fall in the same row of p' are processed in parallel. Using an algorithm similar to [ABF93], witnesses are found for some of these vectors in $O(\log\log m)$ time and $O(m_1 \times m_2)$ work. Two vectors falling in the same row will survive this step only if their difference vector (which is horizontal) has no witnesses in p'.

The actual procedure is described in Appendix I. The basic operation in the algorithm of [ABF93] is a duel between a pair of vectors or a pair of sets of consistent vectors, where vectors v, w are consistent if and only if $w - v$ does not have a witness in p'. Vectors v, w survive this substep only if the horizontal vector $w - v$ does not have a witness in p'.

The witnesses found in this step need not lie completely within p'. We will show that if p' is r-oriented then these witnesses will indeed be within p'; otherwise, for Quad I vectors (Quad II vectors, respectively) these witnesses will be in an $(m'_1 + \lfloor \frac{m'_1}{8} \rfloor) \times (m'_2 + \lfloor \frac{m'_2}{8} \rfloor)$ block with the same top left corner (bottom left corner, respectively) as p'.

We show how each duel is performed so that the witnesses found belong to the regions claimed above. Consider a duel between vectors v, w and suppose that a witness in p' for the vector $w - v$ was computed in Step A. If p' is not r-oriented, since the witness for $w - v$ used in this duel lies in p' and since valid vectors have row length less than $\frac{m'_1}{8}$ and column length less than $\frac{m'_2}{8}$, the witnesses obtained by these duels indeed belong to the region claimed above. Suppose p' is r-oriented. By Lemma 3.1, a new witness $(a, b) \in p'$ for $w - v$ is obtained in constant time and work, where b is in the middle $\frac{m'_2}{4}$ columns of p'. The lengths of v and w imply that either $b + v$ or $b - w$ is in p'. v and w are duelled by comparing b with one of $b + v, b - w$, whichever is in p'. The witness thus obtained is clearly in p'.

Lemma 3.1 *If p' is r-oriented, then given a witness for valid vector v, a witness for v in the middle $\lceil \frac{m'_2}{4} \rceil$ columns of p' (with ties for the center broken arbitrarily) can be found in constant time and work.*

Proof. Let v_1 be a valid horizontal period vector of p'. Recall that $rl(v), rl(v_1) < \frac{m'_1}{8}$ and $cl(v), cl(v_1) < \frac{m'_2}{8}$. Let $a, a+v \in p'$ be the head and tail of some witness for v. There exists an integer i, easily computable in constant time and work, such that $a+iv_1 \equiv a \not\equiv a+v \equiv a+v+iv_1$ and $a + iv_1, a + v + iv_1$ are both in the middle $\lceil \frac{m'_2}{4} \rceil$ columns of p'. $\square$

Substep B.2. All sets of vectors which fall in the same column are processed in parallel in a manner similar to Substep B.1. Two vectors v and w which fall in the same column survive this substep only if the vertical vector $w - v$ does not have a witness in p'.

3.3 Step C

If a surviving valid vector has a witness in p' then a witness for this vector is found in this step. The witness found may not be in p' but will be in p.

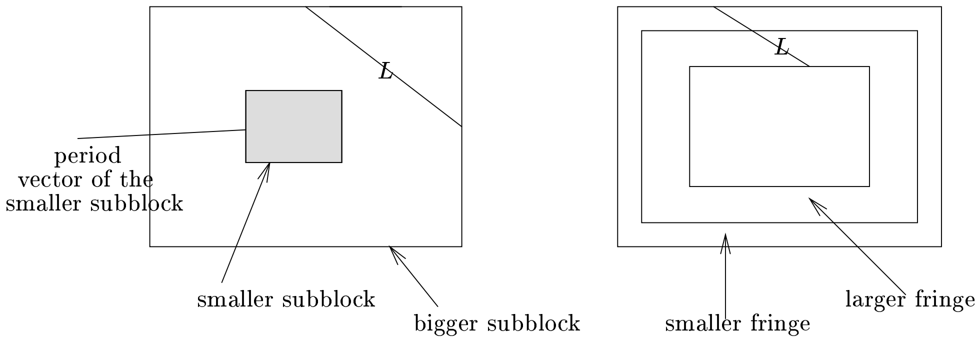

Figure 4: L does not overlap the smaller subblock but will overlap the smaller fringe.

The First Obstacle. Here, we face the first hurdle in generalizing parallel string matching algorithms [BG90] to the 2D case and also in generalizing sequential 2D witness computation algorithms [GP93] to the parallel case.

The obvious generalizations in either case would involve finding witnesses for all non-period vectors of subblocks of p' of successively increasing sizes. The key step would be using witnesses found for a particular subblock to find witnesses for the next largest subblock fast. This step is easily performed in constant time for strings and in time proportional to the size of the larger subblock sequentially.

However, as described in the introduction, it is not obvious how to perform this step in $O(1)$ time and work proportional to the size of the larger subblock. Even in the case when the period vectors of the smaller subblock are all parallel to some line L, it is not clear how this step can be performed in $o(\log\log m)$ time with work linear in the size of the larger subblock. To see why, note that there will be lines parallel to L in the larger subblock which have little or no overlap with the smaller subblock. For such a line L, we will essentially have to solve the string witnesses problem, which takes $\Theta(\log\log|L|)$ time with $O(|L|)$ work.

The Solution: Fringes instead of Subblocks. We take the approach of computing witnesses for non-period vectors of *fringes* of p' of successively increasing sizes. A fringe is illustrated in Fig.4 and Fig.5. Working with fringes ensures that any line in a larger fringe has substantial overlap with the previous (smaller) fringe; this property is critical in obtaining an $O(1)$ time algorithm to find witnesses for all valid non-period vectors of a particular fringe, given the witness for all valid non-period vectors of the previous (smaller) fringe, while maintaining linearity of work. We will outline properties of these fringes and broadly describe the above computation next.

Defining Fringes. We will consider $k = O(\log\log m)$ fringes of p' of successively increasing sizes, the largest (i.e., kth) fringe being p' itself. The *i-fringe* of p', $i = 1, \ldots, k$, is defined to be the set of characters in p' which are less than distance $f_{i,1}$ from the top and bottom boundaries of p' or less than distance $f_{i,2}$ from the left and right boundaries of p' (see Fig.5), for some functions $f_{i,1}$ and $f_{i,2}$ defined as follows.

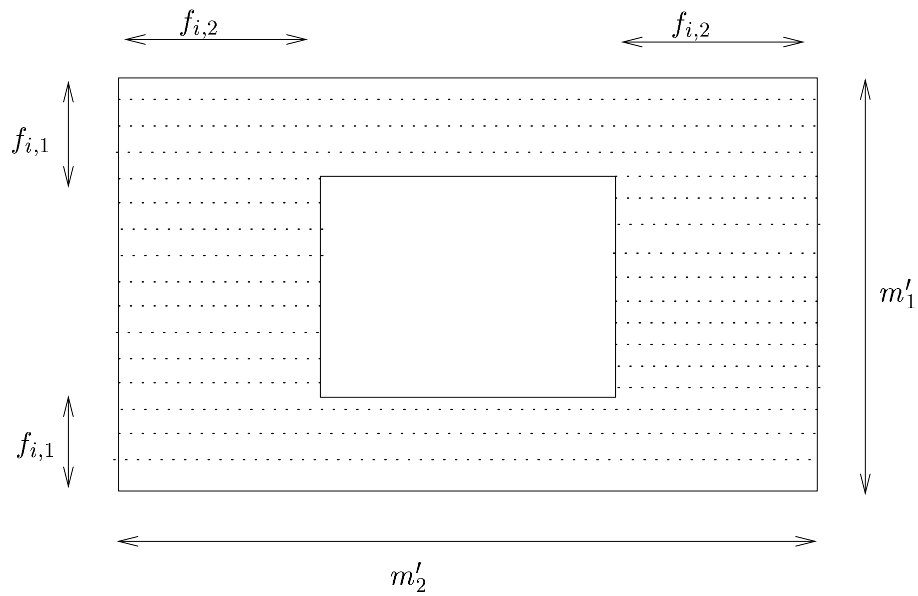

Figure 5: The i-fringe.

Let $i_0 = \lceil \log_{\frac{3}{2}} \log m_1 \rceil$. $f_{0,1}$ is defined to be 1, for convenience.

$$f_{i,1} = \lceil \frac{m_1^{1-(\frac{2}{3})^i}}{\log_2 \log_2 m_1} \rceil, \text{ for } 1 \le i \le i_0.$$
$$f_{i,1} = 2f_{i-1,1}, \text{ for } i = i_0 + 1 \ldots k.$$

k is defined to be the smallest value of i for which $f_{i,1} \ge \frac{m'_1}{2}$; $f_{k,1}$ is now redefined to be $\frac{m'_1}{2}$. It can easily be seen that $k = O(\log \log m)$. $f_{i,2}$ is defined to be $\lceil \frac{m'_2 \times f_{i,1}}{m'_1} \rceil$. Note that since $m_2 \ge m_1 \ge 16$, all terms whose ceilings appear above are greater than 1.

The following properties are satisfied and stated explicitly to make future reference easier.

1. $k \ge 2$, $f_{i,1}, f_{i,2} \ge 2$, and $O(\max\{m_1 \times f_{i,2}, m_2 \times f_{i,1}\}) = O(m_2 \times f_{i,1})$, for all i, $1 \le i \le k$.
2. $\frac{f_{i+1,1}}{f_i,1} \ge 2$, $1 \le i < k-1$.
3. $\frac{f_{i,1}^3}{f_{i-1,1}^2} = O(\frac{m_1}{\log \log m_1})$ for $1 \le i \le i_0$ and $\frac{f_{i,1}^3 f_{i-2,1}^2}{f_{i-1,1}^5} = 2$, for $i_0 + 1 < i \le k-1$.

Computing with Fringes. There are k iterations. These iterations process the fringes in increasing order of size. The following invariants are maintained after iteration i.

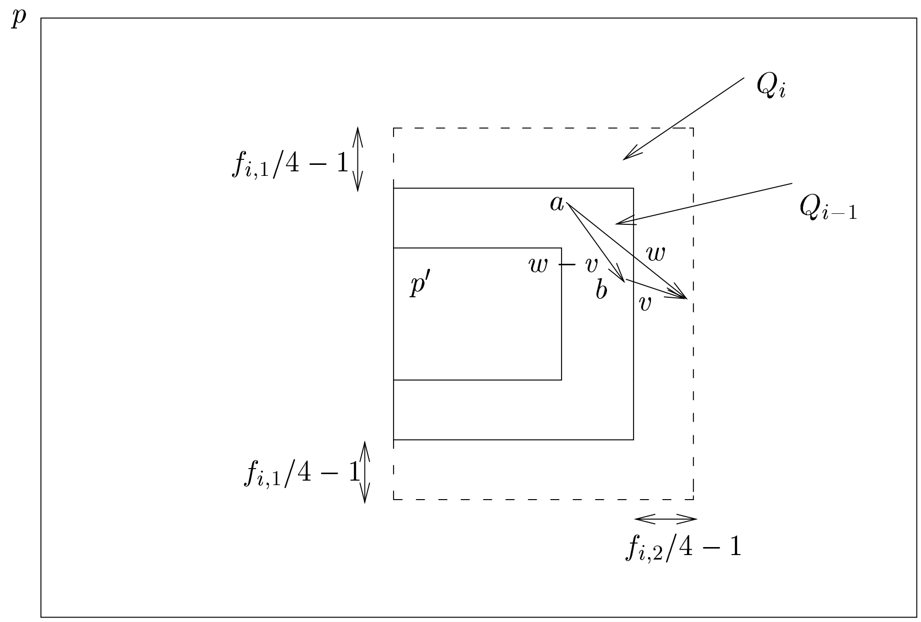

Figure 6: Duel between v and w.

1. Witnesses have been found for all non-period Quad I and Quad II vectors of p' which have row length less than $\frac{f_{i,1}}{4}$ and column length less than $\frac{f_{i,2}}{4}$ and which have a witness in the i-fringe. Clearly, at the end of Step C, witnesses would be found for all valid non-period vectors of p'.

2. If p' is not r-oriented then all witnesses found are contained within either Q_i, where Q_i is a block of size $(m'_1 + 2(f_{k,1}/4 + \sum_{j=1}^{i}(f_{j,1}/4-1))) \times (m'_2 + f_{k,2}/4 + \sum_{j=1}^{i}(f_{j,2}/4-1))$ whose left boundary coincides with that of p' and which extends $f_{k,1}/4 + \sum_{j=1}^{i}(f_{j,1}/4-1)$ rows above and below p' and $m'_2 + f_{k,2}/4 + \sum_{j=1}^{i}(f_{j,2}/4-1)$ columns to the right of p' (see Fig.6).

3. If p' is r-oriented then all witnesses found are inside p'.

Note that the witnesses computed in Steps A and B satisfy the conditions in invariants 2 and 3 with i set to 0.

Witnesses found in Step C indeed lie in p. It is not obvious that Q_k is contained in p. We show that this is indeed the case. Recall that by Property 2, $\frac{f_{i+1,1}}{f_{i,1}} \geq 2$, $1 \leq i < k-1$. Therefore, $\Sigma_{j=1}^{k-1} f_{j,1} \leq 2f_{k-1,1}$ and $\Sigma_{j=1}^{k} f_{j,1} \leq 2f_{k-1,1} + f_{k,1} \leq 3f_{k,1}$. Since $f_{k,1} = \frac{m'_1}{2}$, $f_{k,1}/4 + \sum_{j=1}^{k}(f_{j,1}/4-1) \leq f_{k,1} - k = \frac{m'_1}{2} - k$. Since $k \geq 2$ by Property 1 and $m'_1 \leq \frac{m_1}{2}$, $\frac{m'_1}{2} - k \leq \frac{m_1}{4} - 2 < \lfloor \frac{m_1}{4} \rfloor$, which is the lower bound on the distance between both the upper and the lower boundaries of p and p'. Similarly, $f_{k,2}/4 + \sum_{j=1}^{k}(f_{j,2}/4-1) \leq \frac{1}{4}(\frac{m'_2}{m'_1} f_{k,1} + 1) +$

$\sum_{j=1}^{k}[\frac{1}{4}(\frac{m_2'}{m_1'}f_{j,1}+1)-1] \leq \frac{m_2'}{m_1'}(f_{k,1}/4+\sum_{j=1}^{k}f_{j,1}/4)-k+\frac{(k+1)}{4} \leq \frac{m_2'}{2}-\frac{3k-1}{4} \leq \frac{m_2'}{2}-1 \leq \lfloor\frac{m_2}{4}\rfloor$, which is the lower bound on the distance between the right boundaries of p and p'.

The algorithm for performing the various iterations is described in Section 4. Each iteration takes $O(1)$ time.

Work Done in the various Iterations. The work done in iteration 1 will be $O(\max\{m_1 \times f_{1,2}, m_2 \times f_{1,1}\} \times f_{1,1})$ which equals $O(\frac{m_1 \times m_2}{\log\log m_1})$ by Property 1 and the definition of $f_{1,1}$.

The work done in iteration i, $1 < i \leq k$ will be $O(\max\{m_1 \times f_{i,2}, m_2 \times f_{i,1}\} \times \frac{f_{i,1} \times f_{i,2}}{f_{i-1,1} \times f_{i-1,2}})$, which is $O(m_2 \times \frac{f_{i,1}^3}{f_{i-1,1}^2})$ by Property 1. For $1 < i \leq i_0$, this is $O(\frac{m_2 \times m_1}{\log\log m_1})$ by Property 3 above. The work done in iteration k will be $O(m_2 \times \frac{f_{k,1}^3}{f_{k-1,1}^2}) = O(m_1 \times m_2)$. The work done in iteration $k-1$ will be $O(m_2 \times \frac{f_{k-1,1}^3}{f_{k-2,1}^2}) = O(m_1 \times m_2)$. It remains to consider the work done in iterations $i_0+1, \ldots, k-2$. The ratio of the work done in the ith and $(i-1)$th iterations, $i_0+1 < i \leq k-1$, is $\frac{f_{i,1}^3 f_{i-2,1}^2}{f_{i-1,1}^5}$, which equals 2 by Property 3. Therefore, the total work done in iterations $i_0+1, \ldots, k-1$ is proportional to the work done in iteration $k-1$, which is $O(m_1 \times m_2)$.

3.4 Step D

Recall that vectors surviving Step C are period vectors of p'. This step finds witnesses for those surviving valid vectors which are not period vectors of p. It crucially relies on a new periodicity property which we use to get a $O(\log\log m)$ time and linear work performance for this step. This property is used implicitly in Step D; it is explicitly described in Appendix II.

We outline how Step D is performed for Quad I vectors. An analogous procedure is used to process Quad II vectors. The proofs of the lemmas stated in this section appear in Section 8.

Let S be the set of valid Quad I vectors which still survive. There are three cases.

Case 1. The Singleton Case. Suppose S has at most one vector. This procedure is trivial.

Case 2. The Linear Case. Suppose all vectors in S are parallel. This step is accomplished in $O(\log\log m)$ time and $O(m_1 \times m_2)$ work using $Line(S, p, p)$.

Case 3. The Lattice Case. Suppose S has at least two non-parallel Quad I vectors, v_1 and v_2. Recall that both v_1 and v_2 are period vectors of p'. This case is the hard case and involves exploiting structural properties of the lattice formed by the 2 vectors. We will need the following definitions.

Definitions. The (v_1, v_2)-*lattice points* with respect to a point $a \in p$ are the set of points c such that $c - a$ is a linear combination of v_1 and v_2[1]. A (v_1, v_2)-*lattice path* between two lattice points is a path which consists of consecutive segments, each being one of the vectors $v_1, -v_1, v_2, -v_2$.

Let p'' be the block of size $\frac{m_1'}{2} \times \frac{m_2'}{2}$ concentric with p' (if the center of p'' is not unique then an arbitrary choice is assumed). Consider the (v_1, v_2)-lattice with respect to point e, the top left point of p''. Let C be the lattice cell bounded by the points $e, e+v_1, e+v_2, e+v_1+v_2$.

[1] All references to linear combinations in this paper involve integral coefficients.

The *image* of a point a in p (which may or may not be in C) in cell C is defined to be the character b in C such that $a - b$ is a linear combination of v_1 and v_2 (see Fig.11). a is said to be a *defect* with respect to cell C if a does not match its image in C.

Two Basic Lemmas. The operations in Step D are based on the following two lemmas, whose proofs appear in Section 8. For each Quad I vector v in S, these operations seek to find points x, y such that $y = x + v$ and exactly one of x, y is a defect. If there exists a defect d such that $d + v$ or $d - v$ is in p'', then a witness for p is immediately found.

Lemma 3.2 *p'' contains no defects with respect to C.*

Lemma 3.3 *Let v be a Quad I period vector of p' and let x, y be points in p such that $y = x + v$. If exactly one of x, y is a defect with respect to C then $x \not\equiv y$, i.e., (x, y) is a witness for v. If neither x nor y is a defect with respect to C then $x \equiv y$, i.e., (x, y) is not a witness for v.*

Definitions. Let q be the sub-block of p whose bottom left corner is the bottom left corner of p'' and whose top right corner is the top right corner of p (see Fig.7). Let q' be the sub-block of p whose top right corner is the top right corner of p'' and whose bottom left corner is the bottom left corner of p. Let TR_q be the sub-block of q of size $\lfloor \frac{m_1}{8} \rfloor \times \lfloor \frac{m_2}{8} \rfloor$ with the same top right corner as q. Let $BL_{q'}$ be the sub-block of q' of size $\lfloor \frac{m_1}{8} \rfloor \times \lfloor \frac{m_2}{8} \rfloor$ with the same bottom left corner as q'.

Substeps in Step D. Step D has two main substeps, one which finds witnesses within q and q', and another, which finds witnesses elsewhere in p. This separation is on account of the following complication: for each defect d which is not in q, q' and each vector $v \in S$, either $d - v$ or $d + v$ is in p; this property is not true for defects in q, q'.

Step D.1. For all vectors $v \in S$, a witness in q, if any, is found. Then, for all surviving vectors $v \in S$, a witness in q', if any, is found. Each stage takes $O(\log \log m)$ time and $O(m_1 \times m_2)$ work.

Step D.2. Witnesses in p, if any, are found for all surviving vectors in S. This takes $O(\log \log m)$ time and $O(m_1 \times m_2)$ work.

First, we describe the simpler of the above two steps, Step D.2. If only one vector survives Step D.1 then it is processed as in the Singleton case. If all vectors which survive Step D.1 are parallel then they are processed as in the Linear Case. Otherwise, if at least two non-parallel vectors survive Step D.1, then Claim 1 holds following Step D.1.

Claim 1. If two non-parallel vectors in S survive Step D.1 then all defects in q are in TR_q and all defects in q' are in $BL_{q'}$.

Step D.2. when two non-parallel vectors survive. Suppose that two non-parallel vectors w, w' survive Step D.1. By Claim 1, all defects in q are in TR_q and all defects in q' are in $BL_{q'}$.

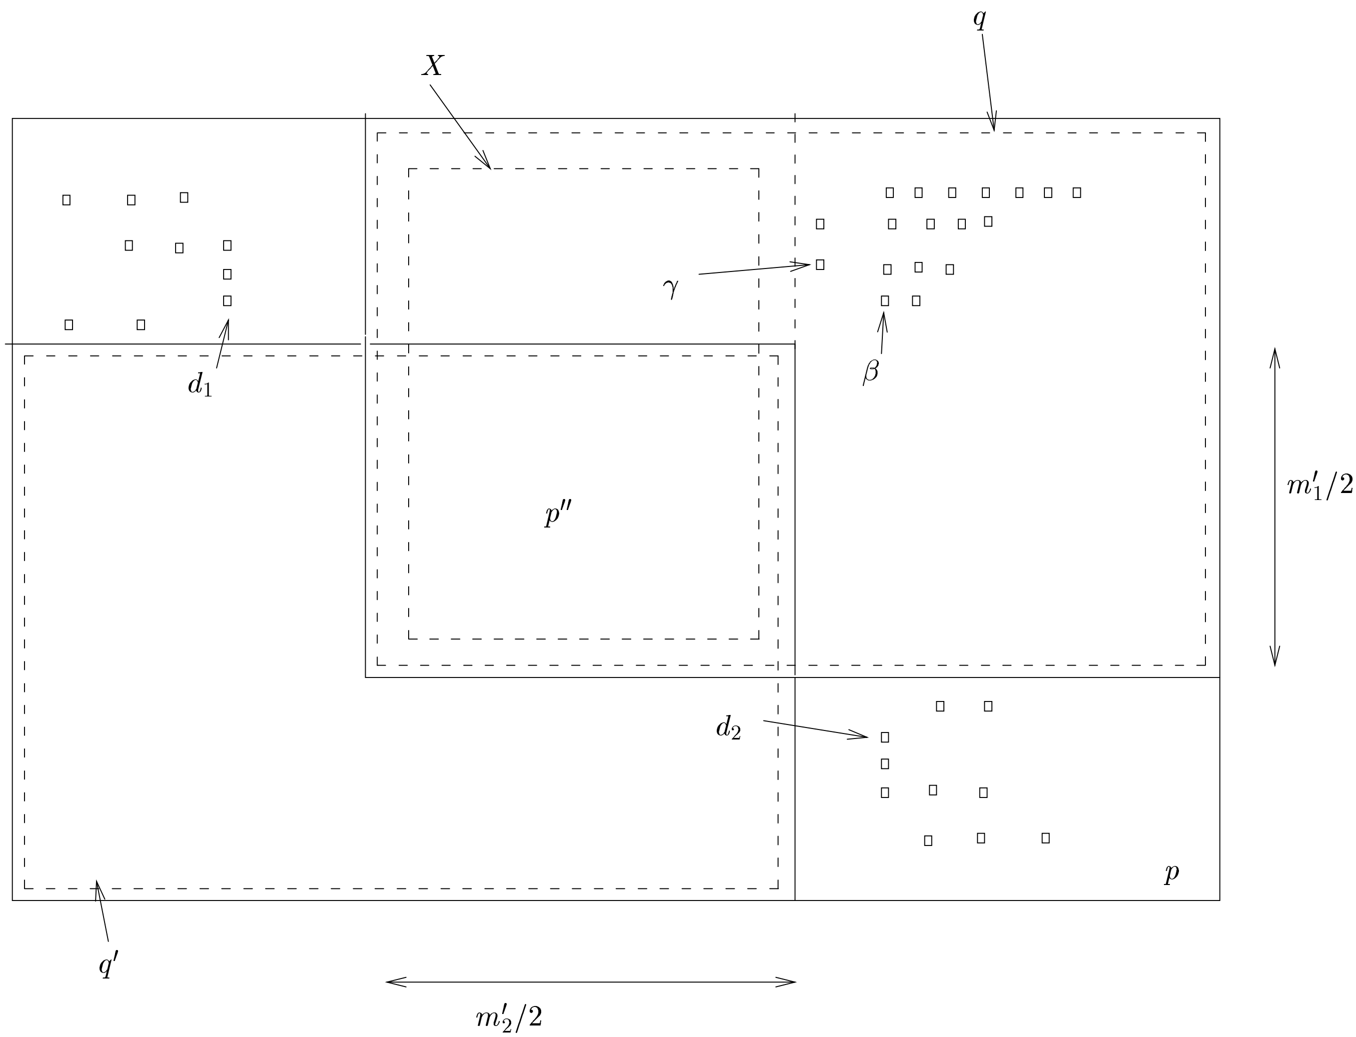

Figure 7: The sub-blocks q and q' with defects.

The bottommost rightmost defect d_1, if any, above q' and to the left of q and the topmost leftmost defect d_2, if any, below q and to the right of q' are found in $O(1)$ time and $O(m_1 \times m_2)$ work. Note that since $rl(v) < \frac{m_1'}{8}$ and $cl(v) < \frac{m_2'}{8}$ for $v \in S$, when the tail of v is on d_1, the head of v can be neither in TR_q nor in $BL_{q'}$; the head of v is therefore on a non-defect. By Lemma 3.3, if d_1 exists then (d_1, d_1+v) is a witness for v. Similarly, if d_2 exists then (d_2, d_2-v) is a witness for v. If neither d_1 nor d_2 exists then by Lemma 3.3, no vector in S which survives Step D.1 has a witness in p.

Step D.1 Outline. Consider only the processing of q. q' is processed similarly. For all but $O(m_1)$ of the vectors in S, it is easy to determine whether or not witnesses exist in q, and to find one, if one exists. This procedure will be described in Section 5.

The Difficult Vectors. Each of the remaining $O(m_1)$ vectors v will have the following property: when the tail of v is placed at the leftmost defect β in the bottommost row of q containing a defect, the head of v is in q and lies on the leftmost defect in its row in q. Finding witnesses, if any, for these vectors in $O(\log \log m)$ time while maintaining $O(m_1 \times m_2)$ work is the second major obstacle in generalizing the sequential 2D witness computation algorithm [GP93].

The above sequential algorithm processed difficult vectors as follows. It constructs a sequence of strings and then finds the leftmost witnesses in each string. The strings themselves are obtained by "peeling off" layers of defects. The first string is derived from the the leftmost defects in each row, the second string from the second leftmost defects and so on. It is not clear how these strings can be obtained in parallel in $O(\log \log m)$ time and $O(m_1 \times m_2)$ work.

Handling Difficult Vectors. We circumvent the problem of peeling layers of defects and computing leftmost witnesses by using a property about the distribution of difficult vectors. We believe that this property is of independent interest. It is described in the following definitions and lemma. Using this property and the procedure *Line*() described in Section 2, witnesses in q, if any, for difficult vectors will be found in $O(\log \log m)$ time and $O(m_1 \times m_2)$ work. This will be described in Section 5.

Definitions. We define the *row partition* of a rectangular sub-block W of p as follows. Let x be the number of rows in W. The rows of W are partitioned into $O(\log x)$ *segments* and *row sets* as follows. The first segment comprises the upper $\lfloor \frac{x}{2} \rfloor$ rows of W; the first row set comprises the remaining rows of W. The jth segment and row set consist of upper $\lfloor \frac{i}{2} \rfloor$ and lower $\lceil \frac{i}{2} \rceil$ rows, respectively, of the $(j-1)$th row set, where i is the number of rows in the $(j-1)$th row set.

Lemma 3.4 *Consider the row partition of Z, the subblock of q comprising the rows above and including the row containing β. If u, v, w are difficult vectors such that rows $rl(u)+1, rl(v)+1, rl(w)+1$ of Z are all in the jth segment and $rl(u) \leq rl(v) \leq rl(w)$, then $v-u$ and $w-v$ are parallel Quad I vectors.*

Further detailed description of Step D.1 is given in Section 5.

4 Step C Description

We need the following primitive before describing the procedure for performing each iteration.

A Useful Primitive. A primitive $Verify(X, Y, V)$ is required in order to describe Step C. Here X and Y are both identically sized rectangular sub-blocks of p or fringes of p' with $X = Y$, and V is a subset of the locations in X with the following property: for all $a, b \in V$, copies of Y placed with the top left corner at a and b match each other wherever they both overlap X. In addition, when X and Y are fringes of p', the region $p' - X$ is strictly below and to the right of all the locations in V (see Fig.16). For each $a \in V$, $Verify(X, Y, V)$ determines a location $bad_a \in X$, if any, such that a copy of Y placed with the top left corner at a mismatches X at location bad_a in X. The procedure for this primitive takes $O(1)$ time and does $O(|X|)$ work. It is described in Appendix I.

The First Iteration. The first iteration is performed as follows. Consider the set S of surviving Quad I vectors with row length less than $\frac{f_{1,1}}{4}$ and column length less than $\frac{f_{1,2}}{4}$. Quad II vectors are processed analogously. Recall from Step B that for any two vectors $u, v \in S$, if $v - u$ is a horizontal vector then no witness exists in p' for $v - u$.

All rows of the 1-fringe X in which vectors in S fall are processed in parallel; clearly there are at most $\frac{f_{1,1}}{4}$ such rows. Consider one such row r and let V be the set of locations in r at which vectors in S fall. Note that by Step B, $b - a$ is a period vector of X for all $a, b \in V$. Therefore copies of X placed with top left corner at a and b match each other. The lengths of the vectors in S imply that all locations in V are strictly above and to the left of the region $p' - X$. $Verify(X, X, V)$ is performed in $O(1)$ time and $O(\max\{m'_1 \times f_{1,2}, m'_2 \times f_{1,1}\})$ work. Suppose $v \in S$ falls on $a \in V$. If bad_a exists then $(bad_a - v, bad_a)$ is a witness for vector v in X and if bad_a does not exist then there is no witness for vector v in X. The work done over all rows in this iteration is $O(\max\{m'_1 \times f_{1,2}, m'_2 \times f_{1,1}\} \times f_{1,1})$ as claimed. Clearly, Invariants 1–3 (described in Section 3.3) hold after the first iteration.

The ith Iteration. We describe the ith iteration assuming that the Invariants 1–3 hold after the $(i-1)$th iteration. We describe only the processing of Quad I vectors; Quad II vectors are processed analogously. The following steps are performed.

Step C.1. Let S be the set of Quad I vectors with row length less than $\frac{f_{i-1,1}}{4}$ and column length less than $\frac{f_{i-1,1}}{4}$ which still survive. Witnesses in the i-fringe, if any, are found for these vectors. This is accomplished in constant time and $O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ work. The actual procedure used to perform this step will be described shortly. Following this step, Invariants 1–3 clearly hold for vectors with row length less than $\frac{f_{i-1,1}}{4}$ and column length less than $\frac{f_{i-1,1}}{4}$.

Step C.2. Let S be the set of surviving Quad I vectors which satisfy at least one of the following two criteria.

1. The row length is at least $\frac{f_{i-1,1}}{4}$ but less than $\frac{f_{i,1}}{4}$ and column length is less than $\frac{f_{i,2}}{4}$.

2. The column length is at least $\frac{f_{i-1,2}}{4}$ but less than $\frac{f_{i,2}}{4}$ and the row length is less than $\frac{f_{i,1}}{4}$.

The $\lfloor \frac{f_{i,1}}{4} \rfloor \times \lfloor \frac{f_{i,2}}{4} \rfloor$ portion of the i-fringe which has the same top left corner as p' is tiled with disjoint smaller blocks of size $\lfloor \frac{f_{i-1,1}}{4} \rfloor \times \lfloor \frac{f_{i-1,2}}{4} \rfloor$ (except at the borders where truncated blocks may be needed).

All smaller blocks are processed in parallel in $O(1)$ time and $O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ work per block; thus the total work done is $O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\} \frac{f_{i,1} \times f_{i,2}}{f_{i-1,1} \times f_{i-1,2}})$. Consider one such block T. Following this step, for every surviving vector v which falls in T, if a witness exists in the i-fringe, then some witness for v satisfying Invariants 2 and 3 will have been found. This is done in two substeps.

Step C.2.1. In this step, witnesses are found for some of the surviving vectors which fall in T. Following this step, all surviving vectors which fall in T are consistent, i.e., a witness for their difference vector does not exist in the i-fringe. This is done in constant time and $O((f_{i-1,1})^2)$ work as follows. Note that $O((f_{i-1,1})^2) = O(m'_2 \times f_{i,1}) = O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ by Property 1 and the fact that $m'_2 \geq m'_1$. There are two cases to consider.

First suppose p' is not r-oriented. Then, by Step B, at most one vector falls in each row of T. Thus there are at most $O(f_{i-1,1})$ vectors which fall in T. All pairs of these vectors are duelled in parallel. Each duel takes constant time and work; the total work done is $O((f_{i-1,1})^2)$ as claimed. Consider a duel between vectors v and w. By Invariant 1 and Step C.1, if a witness exists for $w - v$ in the i-fringe, some witness for $w - v$ would have been found already. By Invariant 2 and Step C.1, such a witness, if any, is in Q_{i-1}. If no witness for $w - v$ has been found, then both v, w survive this duel. Otherwise, a witness is found for at least one of v, w. Since w, v have row lengths less than $\frac{f_{i,1}}{4}$ and column lengths less than $\frac{f_{i,2}}{4}$ and since $w - v$ is neither horizontal nor vertical by Step B, the witness obtained for v or w by this duel is in Q_i (see Fig.6).

Next, suppose p' is r-oriented. Recall that following Step B, all those surviving vectors which fall in the same row of T are consistent, i.e., their difference vector does not have a witness in the whole of p'. By Invariant 3, all witnesses computed so far are within p'. This property plays a key role in this case. The leftmost vector which falls in a row is chosen to be the representative vector for that row and all pairs of representative vectors are duelled in parallel. Each duel takes constant time and work; the total work done is $O((f_{i-1,1})^2)$ as claimed. Each duel is performed in such a way that if a witness is found for a representative vector then a witness is also found for all the surviving vectors which fall in the same row. Consider a duel between representative vectors v and w. If no witness has been found for $w - v$ then by Invariant 1 and Step C.1, there is no witness for $w - v$ in the i-fringe; also, in this case, it is readily seen that there is no witness in the i-fringe for the difference vector of any pair of vectors represented by v and w, respectively. Next, suppose a witness has been found for $w - v$. By Invariant 2, this witness is in p'. Given a witness for $w - v$ in p', by Lemma 3.1, a witness $(a, b) \in p'$ for $w - v$ which lies in the middle $\frac{m'_2}{4}$ columns of p' can be found in constant time and work. The lengths of v, w imply that either $b - w$ or $b + v$ is in p'. Comparing whichever one of these two characters is in p' with b eliminates all vectors which fall in one of the two rows; the resulting witnesses are all within p'.

Step C.2.2. In this step, all surviving vectors which fall in T are either verified to be period vectors of the i-fringe or are eliminated (i.e., a witness is found in the i-fringe). All these vectors are consistent following Step C.2.1, i.e., copies of the i-fringe placed at the characters at which these vectors fall match each other wherever they overlap. Let X denote the i-fringe and let V denote the set of locations in X at which surviving vectors fall in T. The lengths of the vectors which fall in T imply that all locations in V are strictly above and to the left of the region $p' - X$. This step is accomplished in constant time and $O(|X|) = O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ work using $Verify(X, X, V)$. If bad_a exists for some $a \in V$ such that vector v falls at a then $(bad_a - v, bad_a)$ is a witness for v in the i-fringe.

Clearly, at the end of Step C, Invariants 1–3 hold. In the ith iteration, Step C takes $O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\} \frac{f_{i,1} \times f_{i,2}}{f_{i-1,1} \times f_{i-1,2}})$ time, as claimed.

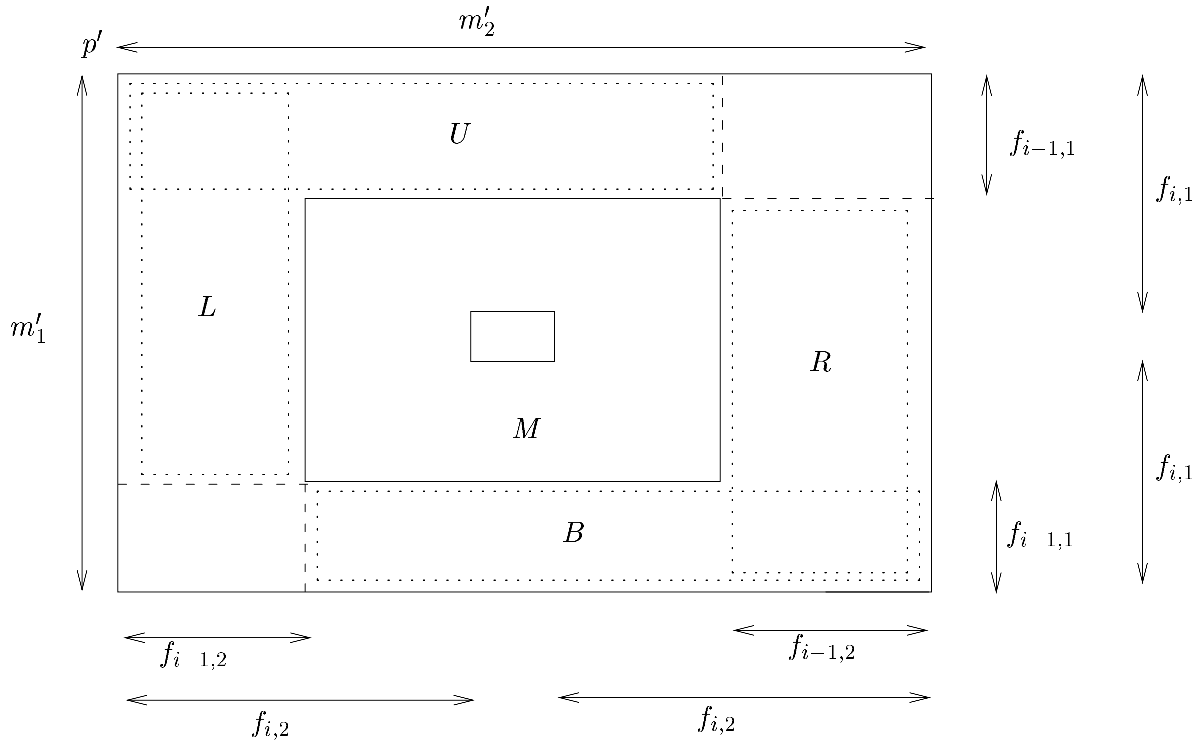

Figure 8: The blocks U, L, R and B.

Step C.1 Description. If S has at most one vector then this procedure is trivial. So assume that S has at least two vectors. Two cases are considered depending upon whether or not the vectors is S are parallel.

Case 1. Suppose all vectors in S are parallel to some line l, say. The i-fringe is partitioned into lines parallel to l. All such lines are considered in parallel. All lines which are contained entirely within the $(i-1)$-fringe contain no witnesses for vectors in S. Without loss of generality, consider a line l' which crosses the top inner boundary of the $(i-1)$-fringe. The lines crossing the other inner boundaries of the $(i-1)$-fringe are handled similarly. For each vector in S, if there exists a witness with tail on l' and head in p', a witness is determined as follows. Let v be the longest vector in S. The uppermost point x in l', if any, such that $x + v \in p'$ and $x \not\equiv x + v$ is found in constant time and $O(|l'|)$ work. The total work done over all lines in the i-fringe is proportional to the size of the i-fringe, i.e., $O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$.

The following lemma holds now. Its proof appears in Section 7.

Lemma 4.1 *If x exists then each vector $w \in S$ has a witness in the i-fringe with head on $x + v$. If x does not exist then there are no witnesses in l' for the vectors in S.*

Case 2. Suppose S has at least two non-parallel Quad I vectors, v_1 and v_2. Four blocks U, L, R, B in the $(i-1)$-fringe are defined as in Fig.8. U and B have size $f_{i-1,1} \times (m'_2 - f_{i-1,2})$

and L and R have size $(m'_1 - f_{i-1,1}) \times f_{i-1,2}$. M denotes the portion of the i-fringe which is not in the $(i-1)$-fringe.

Clearly, if $v \in S$ has a witness in the i-fringe then one or both endpoints of this witness must be in M; in addition, since $rl(v) < \frac{f_{i-1,1}}{4}$, $cl(v) < \frac{f_{i-1,2}}{4}$, and v is a Quad I vector, if exactly one of the endpoints of this witness is in M, the other must be in U, L, B, or R.

Consider the set of (v_1, v_2)-lattice points with respect to e, the top left corner of p'. Let C be the lattice cell bounded by the vertices $e, e+v_1, e+v_2, e+v_1+v_2$. The following lemmas are key. Their proofs follow the description of the procedure which handles this case.

Lemma 4.2 *There are no defects with respect to C in U, L, R or B.*

Lemma 4.3 *Let v be a vector S and let x, y be points in the i-fringe such that $y = x + v$. If exactly one of x and y is a defect with respect to C then $x \not\equiv y$. If x and y are non-defects with respect to C then $x \equiv y$.*

Step C.1 proceeds as follows in this case. The leftmost topmost defect y in either the upper $f_{i,1} - f_{i-1,1}$ rows or the left $f_{i,2} - f_{i-1,2}$ columns of M is found in $O(1)$ time and $O(|M|) = O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ work. Similarly, the rightmost bottommost[2] defect z in either the lower $f_{i,1} - f_{i-1,1}$ rows or the right $f_{i,2} - f_{i-1,2}$ columns of M is found in $O(1)$ time and $O(|M|) = O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$ work.

Next, all vectors $v \in S$ are processed in parallel using constant time and work per vector; the total work done over all vectors is $O(f_{i-1,1} \times f_{i-1,2}) = O(\max\{m'_1 \times f_{i,2}, m'_2 \times f_{i,1}\})$.

Suppose y exists. Since $rl(v) < \frac{f_{i-1,1}}{4}$ and $cl(v) < \frac{f_{i-1,2}}{4}$, $y - v$ is either in M, U, or L. Then by Lemma 4.2 and the definition of y, $y - v$ is not a defect. By Lemma 4.3, $(y - v, y)$ is a witness for v. Similarly, if z exists then $(z, z + v)$ is a witness for v. Next, suppose neither y nor z exists. Then there are no defects in M. By Lemma 4.2, if $x \in M$ and either $x - v$ or $x + v$ is in the i-fringe, then x is a non-defect; by Lemma 4.3, there is no witness for v with one endpoint in M. Therefore, there is no witness for v in the i-fringe.

Proving Lemmas 4.2 and 4.3. will require the machinery developed in Section 6. The proofs are given in Section 7.

5 Step D.1 Description

Recall that we need to describe only the processing of q, i.e., finding witnesses in q, if any, for vectors in S. Here S is the set of valid vectors which survive at the beginning of Step D.1. From Step C, these vectors do not have witnesses in p' and therefore in p''. We need the following definition.

Definition. A b-defect in a subblock A is the bottommost defect in its column and an l-defect is the leftmost defect in its row. A vector v is said to be *l-consistent* in a subblock A if for every l-defect $a \in A$, each of $a + v$ and $a - v$ is either outside A or also an l-defect. *r-consistency*, *t-consistency* and *b-consistency* are defined analogously.

If q has no defects then, by Lemma 3.3, all vectors in S survive this step. So assume that q has a defect. The b-defects in each column of q and the l-defects in each row of q are found

[2] The rightmost defect in the bottommost row containing a defect.

in $O(1)$ time and $O(m_1 \times m_2)$ work. There are three further substeps, Steps D.1a, D.1b and D.1c, each of which takes $O(\log \log m)$ time and $O(m_1 \times m_2)$ work.

Step D.1a. Witnesses will be determined for all but two sets of vectors in S, denoted S' and S''. There will be $O(m_1)$ vectors in S'; these vectors are the difficult vectors and may or may not have witnesses in q. Vectors in S'' will not have witnesses in q. Further, if S'' is non-empty, all defects in q will be in the top right corner sub-block of q of size $\lfloor \frac{m_1'}{8} \rfloor \times \lfloor \frac{m_2'}{8} \rfloor$. This step takes $O(\log \log m)$ time and $O(m_1 \times m_2)$ work and proceeds as follows.

See Fig.7. Let γ be the bottommost leftmost defect in q. Let β be the leftmost bottommost defect in q. γ and β along with all the l-defects and b-defects in q are found in constant time and $O(m_1 \times m_2)$ work.

There are two cases to consider. In each case, each surviving vector in S is processed in constant time and work for a total of $O(m_1 \times m_2)$ work, plus, in Case 1, an additional $O(\log \log m)$ time and $O(m_1 \times m_2)$ work is used.

Case 1. β is in the left $\lfloor \frac{m_2'}{8} \rfloor$ columns of q. By Lemma 3.2, p'' has no defects and therefore, β is above p''. For all surviving non-horizontal vectors $v \in S$, $\beta + v \in q$ and $(\beta, \beta + v)$ is a witness for v by Lemma 3.3. All surviving horizontal vectors $v \in S$ are then processed as in the Linear Case (i.e., the case in which all vectors in S were parallel). S' and S'' are both empty in this case.

Case 2. β is not in the left $\lfloor \frac{m_2'}{8} \rfloor$ columns of q.

Let *br* be the index of the row containing β. In this case, surviving vectors in S with row length less than *br* are processed separately from those whose row length is at least *br*.

For each surviving vector $v \in S$ with row length less than *br*, $\beta - v \in q$. Witnesses are sought for each such vector v in one of the following two positions: with head at β or with tail at the l-defect in q in the row containing $\beta - v$. Clearly, one of these two positions provides a witness for v by Lemma 3.3 unless $\beta - v$ is the l-defect in its row. Those vectors v for which $\beta - v$ is an l-defect in q comprise the set S'. S' has at most one vector for each row length and therefore $|S'| = O(m_1)$.

For each surviving vector $v \in S$ with row length at least *br*, a witness is sought with tail at γ. Clearly, if such a vector v exists then β and hence all defects in q are in the upper $br < \frac{m_1'}{8}$ rows of q. If $\gamma + v \in q$ then $\gamma + v$ is below β and by Lemma 3.3, $(\gamma, \gamma + v)$ is a witness for v. Otherwise, if $\gamma + v \notin q$ then $\gamma + v$ is to the right of the right boundary of q and $\beta - v$ is above the upper boundary of q. Then, for all defects $\delta \in q$, $\delta + v$ and $\delta - v$ are both outside q and, by Lemma 3.3, there are no witnesses for v in q; these vectors v comprise set S''. Clearly, if set S'' is non-empty then all defects in q are in the top right corner sub-block of q of size $\lfloor \frac{m_1'}{8} \rfloor \times \lfloor \frac{m_2'}{8} \rfloor$.

Step D.1b. This step eliminates those vectors in S' which have a witness in q with one endpoint on either an l-defect or a b-defect in q.

All vectors in S' are processed in parallel in this step. Consider vector $v \in S'$. For every l-defect and b-defect α in q, α is compared to $\alpha - v$ and $\alpha + v$. The time taken by this step in $O(1)$; the work done is $O(m_1 \times m_2)$, $O(m_2)$ per vector in S'.

Definition. Let X be the sub-block of q comprising its left $\frac{m_2'}{2}$ columns.

Lemma 5.1 and Lemma 5.2 hold following this step. Their proofs along with the proofs of the other lemmas in this section are given in Section 8. Claim 1 of Step D is then satisfied for q by Corollary 5.3; it can be shown for q' in a similar manner.

Lemma 5.1 *Consider a vector v in either S' or S'' which survives Steps survives Steps D.1a and D.1b. If γ is not in the left $\lfloor \frac{m'_2}{8} \rfloor$ columns of q, v is l-consistent in q. If γ is in the left $\lfloor \frac{m'_2}{8} \rfloor$ columns of q, v is b-consistent in X.*

Lemma 5.2 *Let v, w be non-parallel valid Quad I vectors. Let Y be any sub-block of q containing at least the $\frac{m'_1}{2}$ bottommost rows of q. If v, w are l-consistent in Y then all defects in Y are in the top right corner sub-block Y' of Y of size $(rl(v)+rl(w)-1) \times (cl(v)+cl(w)-1)$. Further, if γ in q is in the left $\lfloor \frac{m'_2}{8} \rfloor$ columns of q then v and w cannot both be b-consistent in X.*

Corollary 5.3 *If two non-parallel vectors $v, w \in S$ survive Steps D.1a and D.1b then v, w are l-consistent in q and all defects in q are in the top right corner sub-block of q of size $(rl(v)+rl(w)-1) \times (cl(v)+cl(w)-1)$. Further this sub-block is contained in TR_q.*

Proof. γ is outside the left $\lfloor \frac{m'_2}{8} \rfloor$ columns of q; otherwise, by Lemma 5.1, v and w are b-consistent in X which contradicts Lemma 5.2. By Lemma 5.1, v and w are l-consistent in q. The rest of the corollary follows from Lemma 5.2 (with q replacing Y) and the fact that $rl(v)+rl(w) < 2\frac{m'_1}{8} \leq \frac{m_1}{8}$ and $cl(v)+cl(w) < 2\frac{m'_2}{8} \leq \frac{m_2}{8}$. $\square$

Lemma 5.4 *Suppose two non-parallel vectors $v, w \in S'$ survive Steps D.1a and D.1b, where $rl(w) \geq rl(v)$. Then $w - v$ is a Quad I vector.*

Step D.1c. This step is performed only if at least two non-parallel vectors in S' survive Step D.1b. By Corollary 5.3, all surviving vectors $v \in S'$ are l-consistent in q and all defects in q are in TR_q. Witnesses in q are found for all surviving vectors in S' in this step in $O(\log\log m)$ time and $O(m_1 \times m_2)$ work.

Let br be the index of the bottommost row in q containing a defect. Let Z be the sub-block of q comprising the top br rows of q. We now show that any witness in q for a surviving vector $v \in S'$ must have both endpoints in Z. By Lemma 3.3, one endpoint of any witness for v must be in Z. Suppose for a contradiction that $(a, a+v)$ is a witness for v where $a \in Z, a+v \notin Z$. Then, if b is the l-defect in the row containing a then $b+v$ is in q; further, it is not a defect. But, by Corollary 5.3, v is l-consistent in q, which is a contradiction.

Consider the row partition of Z. All segments are processed in parallel. Consider the jth segment. Let R denote the sub-block comprising the rows of the jth segment and the jth row set, i.e., R is the $j-1$th row set if $j > 1$ and $R = Z$ if $j = 1$. Let V be the set of surviving vectors v in S' such that row $rl(v)+1$ of Z is in the jth segment. The difference vectors of all pairs of vectors in V are parallel by Lemma 3.4. Any witness in Z for a vector in V must have its head in R. Witnesses with heads in R and tails in Z, if any, for vectors in V are found in $O(\log\log m)$ time and $O(|R|) = O(\frac{br}{2^{j-1}} \times m_2)$ work using $Line(V, R, Z)$. Thus, the total work done over all segments is $O(m_1 \times m_2)$.

6 Some Useful Lemmas

Let A be a sub-block of p with r_A rows and c_A columns. Imagine A to be placed on an infinite grid so that each location in A coincides with some grid point. Henceforth, we refer to these grid points simply as *points*. Let v_1 and v_2 be two non-parallel vectors.

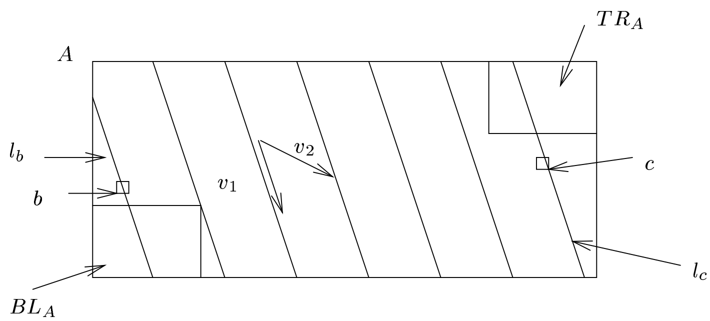

Figure 9: The set of lines L.

Definitions. We define the set $good_A(v_1, v_2) \subseteq A$ as follows. Let BL_A, BR_A, TL_A, TR_A be, respectively, the bottom left, bottom right, top left and top right corner sub-blocks of A of size $(rl(v_1)+rl(v_2)-1) \times (cl(v_1)+cl(v_2)-1)$ (i.e., number of rows times number of columns). If v_1 and v_2 are Quad I vectors then $good_A(v_1, v_2) = A - BL_A - TR_A$. If v_1 and v_2 are Quad II vectors then $good_A(v_1, v_2) = A - BR_A - TL_A$. If v_1 is a Quad I vector and v_2 is a Quad II vector or if v_1 is a Quad II vector and v_2 is a Quad I vector then $good_A(v_1, v_2) = A$.

Lemma 6.1 *Suppose $r_A \geq rl(v_1) + rl(v_2)$ and $c_A \geq cl(v_1) + cl(v_2)$. Let $b, c \in A$ be (v_1, v_2)-lattice points with respect to some point $a \in A$. If $b, c \in good_A(v_1, v_2)$ and v_1 and v_2 are both Quad I or Quad II vectors then there is a (v_1, v_2)-lattice path completely contained in A between points b and c.*

Proof. Without loss of generality, assume that v_1 and v_2 are both Quad I vectors and that v_1 is clockwise with respect to v_2.

Let l_b, l_c be the lines in A, parallel to v_1, through b and c, respectively, Consider the set L of lines in A parallel to v_1, at separation v_2, which lie between l_b and l_c, both lines inclusive (see Fig.9). Without loss of generality, assume that l_b is to the left of l_c. We show that for each pair of adjacent lines $l_1, l_2 \in L$, there is a lattice point $e \in l_1$ with $e + v_2$ in A and hence in l_2. It follows that there is a lattice path from b to c.

Since $b, c \in good_A(v_1, v_2)$, l_b and l_c both span either at least $rl(v_1) + rl(v_2)$ rows or at least $cl(v_1) + cl(v_2)$ columns. Without loss of generality, assume that each spans at least $rl(v_1) + rl(v_2)$ rows. Since l_1, l_2 are parallel to and between l_b and l_c, l_1 and l_2 each span at least $rl(v_1) + rl(v_2)$ rows. It follows that both l_1 and l_2 have at least one lattice point.

Suppose, for a contradiction that for all lattice points $e \in l_1$, $e + v_2 \notin A$.

The portion of l_1 in the top $rl(v_1)$ rows spanned by it must contain a lattice point. Let f be the topmost lattice point in l_1, i.e., $f - v_1 \notin A$. Note that since l_1 spans at least $rl(v_1) + rl(v_2)$ rows, $f + v_2$ cannot be below the lowest of the rows spanned by l_1. Then, by the assumption that $f + v_2 \notin A$, $f + v_2$ must be to the right of the right boundary of A. Since $c_A \geq cl(v_1) + cl(v_2)$, $f - v_1$ cannot be to the left of the left boundary of A. Since $f - v_1 \notin A$, $f - v_1$ is above the top boundary of A. It follows that f is in the top $rl(v_1)$ rows and right $cl(v_2)$ columns of A. Since v_1 is clockwise with respect to v_2, the only rows which l_2 can then span are the rows above $f + v_2$ (see Fig.10). But there are only $rl(v_1) + rl(v_2) - 1$ such rows, a contradiction. □

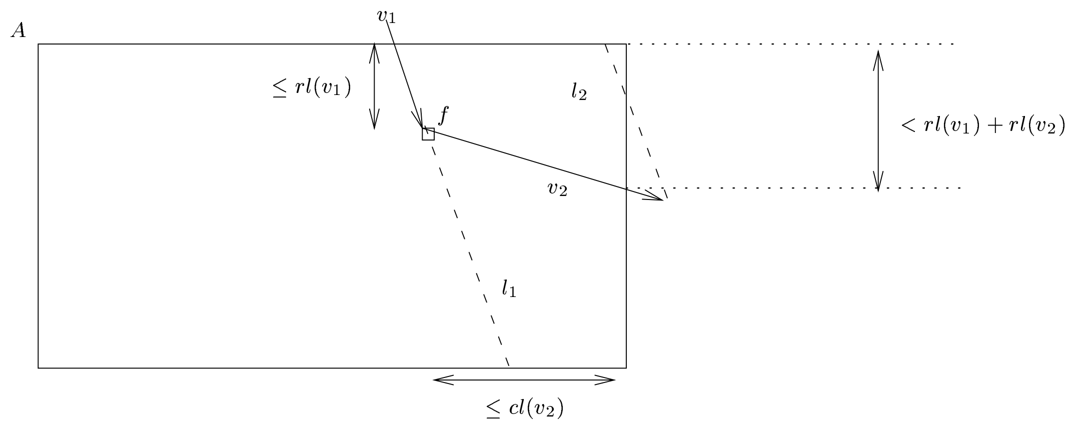

Figure 10: Rows possibly spanned by l_2.

Lemma 6.2 *Suppose $r_A \geq rl(v_1) + rl(v_2)$ and $c_A \geq cl(v_1) + cl(v_2)$. Let $b, c \in A$ be (v_1, v_2)-lattice points with respect to some point $a \in A$. If v_1 is a Quad I vector and v_2 is a Quad II vector then there is a (v_1, v_2)-lattice path completely contained in A between points b and c.*

Proof. Let l_b, l_c, l_1, l_2 be as in the proof of Lemma 6.1. We show that there is a lattice point $e \in l_1$ with $e + v_2$ in A and hence in l_2. It follows that there is a lattice path from b to c.

First, we show that l_1 and l_2 have at least one lattice point each. Suppose for a contradiction that l_1 does not have any lattice points. Then l_1 spans fewer than $rl(v_1)$ rows and $cl(v_1)$ columns. Therefore, l_1 intersects either the top and right boundaries of A or the lower and left boundaries of A. Suppose l_1 intersects the top and right boundaries of A; the other case is handled similarly. Then l_c also intersects the top and right boundaries of A and the line parallel to $-v_2$ through c must intersect l_1. This contradicts the fact that l_1 has no lattice points. l_2 is shown to have a lattice point in the same way.

Next, suppose for a contradiction that for all lattice points $e \in l_1$, $e + v_2 \notin A$.

Let g be any lattice point on l_2 and let f be the lattice point on l_1 closest to $g - v_2$. Let l be the line obtained by extending l_1 indefinitely on both sides. Clearly, f, $f - v_1$ and $g - v_2$ are lattice points on l. If $f = g - v_2$, a contradiction results. So assume that f is below $g - v_2$ on l; the case when it is above $g - v_2$ is handled similarly. Then since $f - v_1$ and f are consecutive lattice points on l and $g - v_2$ is also a lattice point on l, $g - v_2$ either coincides with or is above $f - v_1$ in l.

If $g - v_2 \in A$ then there is a lattice point $g - v_2 \in l_1$ such that $g - v_2 + v_2 = g \in A$, a contradiction. Similarly, if $f + v_2 \in A$, a contradiction results. So suppose that $g - v_2, f + v_2 \notin A$. Since f is below $g - v_2$ and f is the lattice point on l_1 closest to $g - v_2$, $f - v_1 \notin A$. Since $c_A \geq cl(v_1) + cl(v_2)$, either $f + v_2$ is to the right of the right boundary of A and $f - v_1$ is above the top boundary of A or $f + v_2$ is above the top boundary of A and $f - v_1$ is to the left of the left boundary of A or both $f + v_2, f - v_1$ are above the top boundary of A. Recall from the previous paragraph that $g - v_2$ must is aligned with or above $f - v_1$. Since v_2 is a Quad II vector, g itself must be aligned with or above $f - v_1$. Since f is below $g - v_2$, g is aligned with or above $f + v_2$. It follows that in all the above three cases, g is above the top boundary of A, a contradiction. $\Box$

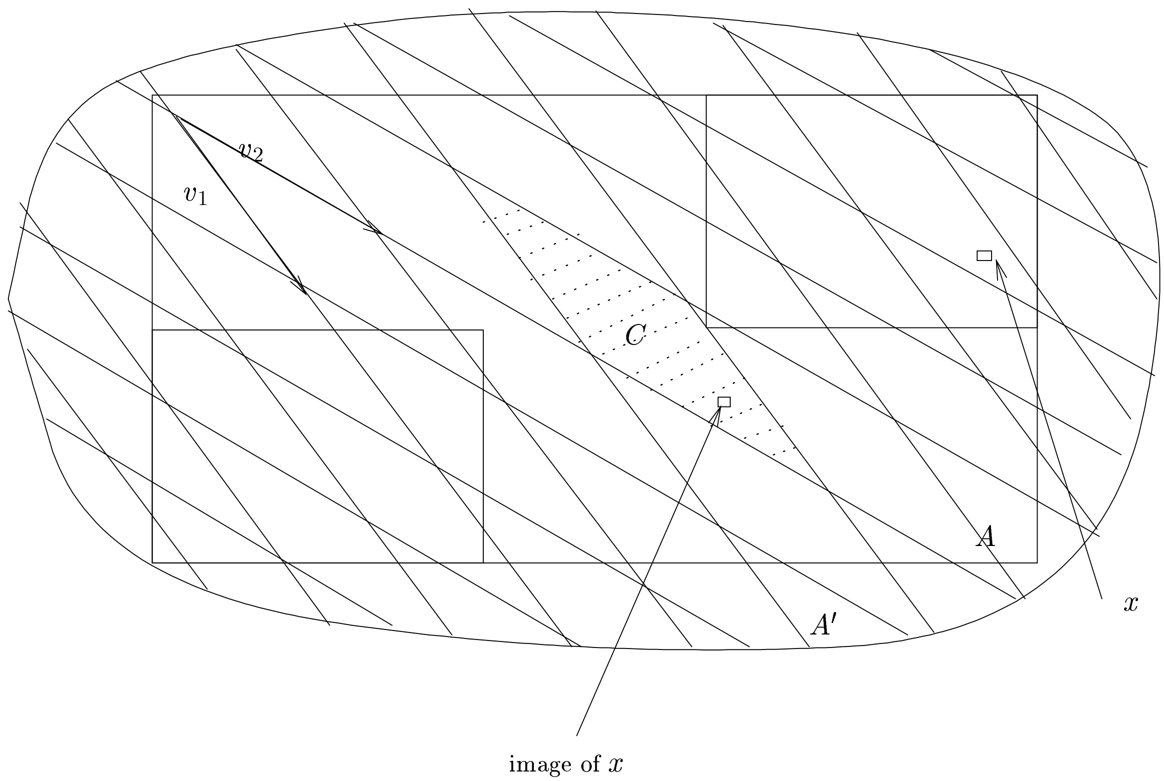

Figure 11: Cell C of Lemma 4.3 and the image of x in C. $good_A(v_1, v_2)$ is the area in A outside the two subrectangles in A.

Remark. Note that Lemmas 6.1 and 6.2 also appear implicitly in [GP92, GP93].

Lemma 6.3 *Suppose $r_A \geq rl(v_1)+rl(v_2)$ and $c_A \geq cl(v_1)+cl(v_2)$. Further, suppose v_1, v_2 are period vectors of A. Let C be any cell in the (v_1, v_2)-lattice with respect to some point in A' with the property that all points in C are also in $good_A(v_1, v_2)$ (see Fig.11). If $x \in good_A(v_1, v_2)$ then x cannot be a defect with respect to C.*

Proof. Let x' be the image of x in C. Consider the (v_1, v_2)-lattice with respect to x'. Since x and x' are lattice points in this lattice and since $x, x' \in good_A(v_1, v_2)$, by Lemmas 6.1 and 6.2, there is a (v_1, v_2)-lattice path in A from x to x'. Since v_1, v_2 are period vectors of A, $x \equiv x'$. □

Definitions. An *l-defect* in A with respect to cell C is the leftmost defect with respect to C in its row. An *r-defect* in A with respect to cell C is the rightmost defect with respect to C in its row. Similarly, a *t-defect* is the topmost defect in its column and a *b-defect* is the bottommost defect in its column.

A vector v is said to be *l-consistent* in A if for every l-defect $a \in A$, each of $a + v$ and $a - v$ is either outside A or also an l-defect. *r-consistency*, *t-consistency* and *b-consistency* are defined analogously.

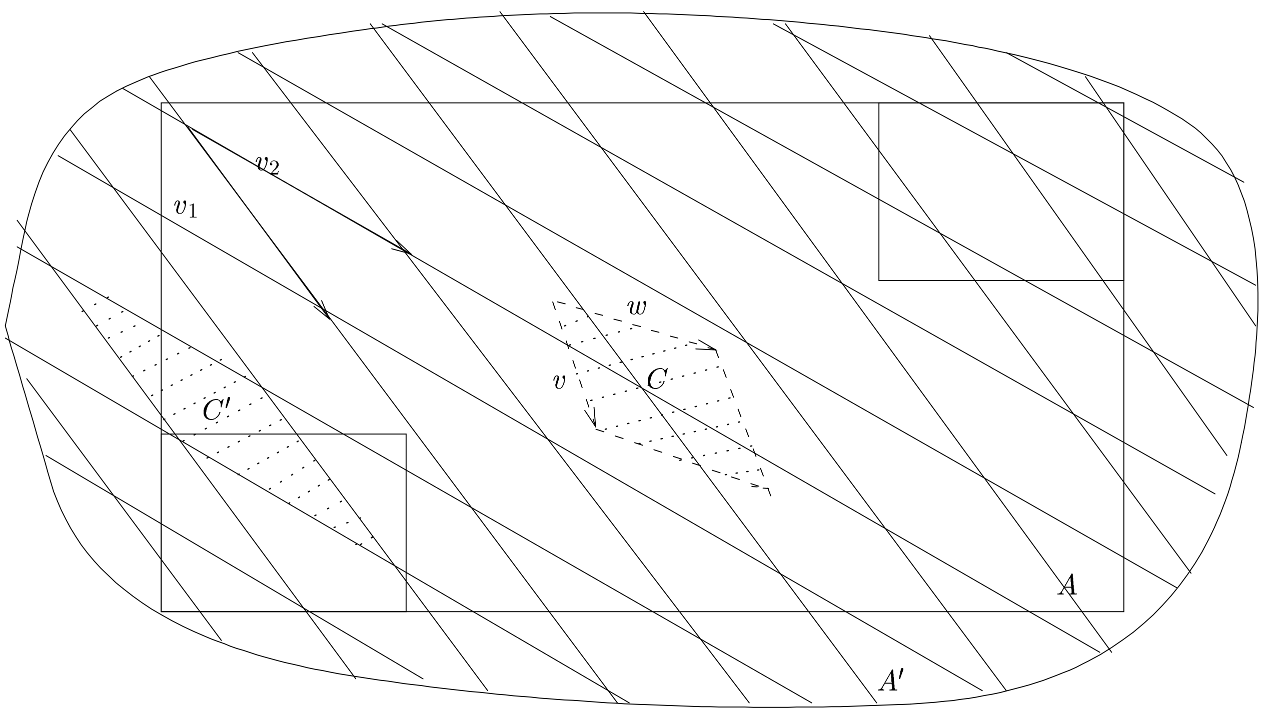

Figure 12: The cells C, C' in Lemma 4.4. $good_A(v, w)$ is the area in A outside the two subrectangles in A.

Lemma 6.4 *Let v, w be two non-parallel vectors. Suppose $r_A \geq rl(v) + rl(w)$ and $c_A \geq cl(v) + cl(w)$. Let C' be a cell in the (v_1, v_2)-lattice with respect to some point in A', where v_1 and v_2 are two non-parallel vectors (see Fig.12). Suppose there exists a cell C in the (v, w)-lattice with respect to some point in A' with the property that all points in C are in $good_A(v, w)$ and none is a defect with respect to cell C'. If v, w are l-consistent (r-consistent, b-consistent, t-consistent, respectively) in A then all l-defects (r-defects, b-defects, t-defects, respectively) in A with respect to C' are outside $good_A(v, w)$.*

Proof. Consider the case when v, w are l-consistent; the other cases are similar. Suppose there is an l-defect a in A with respect to C' such that $a \in good_A(v, w)$. Then by Lemmas 6.1 and 6.2, there is a (v, w)-lattice path in A from a to a', its image in C. Since v and w are l-consistent, every point on this path must be an l-defect with respect to C'. But a' is not a defect, a contradiction. □

Definition. A cell C in the (v_1, v_2)-lattice with respect to some point in A is said to be *safe* for vector v in A if for all $b \in C$, $b + v$ is in A and $b + v$ is not a defect with respect to C.

Lemma 6.5 *Let v be a period vector of A and let x, y be points in A' such that $y = x + v$. Suppose there exists a cell C in A in the (v_1, v_2)-lattice with respect to some point in A with the property that C is safe for vector v in A. If x' and y' are the images in C of x and y, respectively, then $x' \equiv y'$.*

Proof. Let $z = x' + v$. Since C is safe for v in A, $z \in A$ and z matches its image in C. Since v is a period vector of A, $z \equiv x'$. It then suffices to show that y' is the image of z with respect to C, and therefore $y' \equiv z \equiv x'$.

Note that $x' - x = (x' + v) - (x + v) = z - y$. Since $x' - x$ and $y' - y$ are linear combinations of v_1 and v_2, $(y' - y) - (x' - x) = y' - z$ is a linear combinations of v_1 and v_2. □

Corollary 6.6 *Let v be a period vector of A and let x, y be points in A' such that $y = x + v$. Suppose there exists a cell C in the (v_1, v_2)-lattice with respect to some point in A with the property that all points in C are in A and C is safe for vector v in A. If exactly one of x, y is a defect with respect to C, then $x \not\equiv y$, i.e., (x, y) is a witness for v. If neither x nor y is a defect with respect to C, then $x \equiv y$, i.e., (x, y) is not a witness for v.*

Proof. Let x' and y' be the images of x and y, respectively, in C. By Lemma 6.5, $x' \equiv y'$. If exactly one of x, y is a defect with respect to C then exactly one of the equalities $x \equiv x', y \equiv y'$ is true and therefore either $x \not\equiv x' \equiv y' \equiv y$ or $x \equiv x' \equiv y' \not\equiv y$. If neither x nor y is a defect with respect to C then $x \equiv x' \equiv y' \equiv y$. □

7 Step C Proofs

Proof of Lemma 4.1. First, suppose x exists. We show that $x + v \not\equiv x + v - w$. Note that $x \not\equiv x + v$, v is a Quad I vector, and that l' crosses the top inner boundary of the $(i-1)$-fringe. $x + v$ cannot lie in the $(i-1)$-fringe, otherwise both $x + v$ and x are in the $i - 1$-fringe; this contradicts the fact that there is no witness for v in the $(i-1)$-fringe. Since $rl(v), rl(w) < \frac{f_{i-1,1}}{4}$ and $cl(v), cl(w) < \frac{f_{i-1,2}}{4}$, there exists a number j such that $x + v - jv$ and $x + v - w - jv$ are both in the $(i-1)$-fringe. Since there is no witness for w in the $(i-1)$-fringe, $x + v - jv \equiv x + v - w - jv$. Since x is the uppermost point l' such that $x + v \in p'$ and

$x \not\equiv x+v$, $x+v \not\equiv x+v-v \equiv x+v-jv$ and $x+v-w \equiv x+v-w-jv$, and therefore, $x+v \not\equiv x+v-w$.

Next, suppose x does not exist. We show that for every $w \in S$ and $y \in l'$ such that $y+w \in p'$, $y \equiv y+w$. This is true if both y and $y+w$ are in the $(i-1)$-fringe as there is no witness for w in the $(i-1)$-fringe. So suppose at least one of $y, y+w$ is not in the $(i-1)$-fringe. Since $rl(v), rl(w) < \frac{f_{i-1,1}}{4}$ and $cl(v), cl(w) < \frac{f_{i-1,2}}{4}$, there exists a number j such that $y+w-jv$ and $y-jv$ are both in the $(i-1)$-fringe. Since there is no witness for w in the $(i-1)$-fringe, $y+w-jv \equiv y-jv$. Since v is the longest vector in S, $y+w-v \in l'$. Then, by the assumption that x does not exist, $y+w \equiv y+w-v \equiv y+w-jv$ and $y \equiv y-jv$, and therefore, $y+w \equiv y$, as claimed. □

Proof of Lemma 4.2: Let U', L', R', B' be the sub-blocks of p' comprising the upper $f_{i-1,1}$ rows, left $f_{i-1,2}$ columns, right $f_{i-1,2}$ columns and bottom $f_{i-1,1}$ rows of the $(i-1)$-fringe, respectively. Let $BL_{U'}$ and $TR_{U'}$ be, respectively, the bottom left and top right corner sub-blocks of U' of size $\lfloor\frac{f_{i-1,1}}{2}\rfloor \times \lfloor\frac{f_{i-1,2}}{2}\rfloor$. $BL_{L'}$, $TR_{L'}$, $BL_{R'}$, $TR_{R'}$, $BL_{B'}$ and $TR_{B'}$ are defined analogously.

Note that $rl(v_1), rl(v_2) < \frac{f_{i-1,1}}{4}$, $cl(v_1), cl(v_1) < \frac{f_{i-1,2}}{4}$, and that U' and L' have sizes $f_{i-1,1} \times m'_2$ and $m'_1 \times f_{i-1,2}$, respectively, where $m'_1 \geq 2f_{i-1,1}$ and $m'_2 \geq 2f_{i-1,2}$.

Consider U' first. All points in U' except those in $TR_{U'}$ or $BL_{U'}$ are in $good_{U'}(v_1, v_2)$. In addition, C does not overlap either of these blocks; therefore all points in C are also in $good_{U'}(v_1, v_2)$. By Lemma 6.3 (with A replaced by U'), all points in U' except those in $BL_{U'}$ and $TR_{U'}$ are non-defects with respect to C.

Similarly, all points in L' except those in $BL_{L'}$ and $TR_{L'}$ are non-defects with respect to C.

Note that $TR_{U'}, BL_{L'}$ are outside U and L. Further, $TR_{L'} \subset U' - BL_{U'} - TR_{U'}$ and therefore all points in $TR_{L'}$ are non-defects with respect to C. Similarly, $BL_{U'} \subset L' - TR_{L'} - BL_{L'}$ and therefore all points in $BL_{U'}$ are non-defects with respect to C. It follows that there are no defects with respect to C in either U or L.

Next, we show that there are no defects in R or in B. First, we claim that for all $a \in good_{R'}(v_1, v_2)$, a is not a defect with respect to C. Then, since $TR_{R'}$ is outside R, the only defects with respect to C in R are in $BL_{R'}$. Similarly, the only defects with respect to C in B are in $TR_{B'}$. Since $BL_{R'}$ and $TR_{B'}$ do not overlap, there are no defects with respect to C either in R or in B.

Let $a \in good_{R'}(v_1, v_2)$ and a' be the image of a in C, i.e., a is a (v_1, v_2)-lattice point with respect to a'. From the sizes of v_1, v_2, it follows that there is a (v_1, v_2)-lattice point b with respect to a' in $R' \cap U' - TR_{R'}$. Clearly, $b \in good_{R'}(v_1, v_2)$ and $good_{U'}(v_1, v_2)$. By Lemma 6.1 applied to R' and then to U', there is a (v_1, v_2)-lattice path in R' from a to b and another in U' from b to a'. Since v_1 and v_2 are period vectors of the $(i-1)$-fringe, $a \equiv a'$ and a is not a defect, as claimed. □

Proof of Lemma 4.3: Recall that $rl(v), rl(v_1), rl(v_2) < \frac{f_{i-1,1}}{4}$ and $cl(v), cl(v_1), cl(v_2) < \frac{f_{i-1,2}}{4}$, and that U has size $f_{i-1,1} \times (m'_2 - f_{i-1,2})$. Further $m'_1 \geq 2f_{i-1,1}$ and $m'_2 \geq 2f_{i-1,2}$.

By Lemma 4.2, there are no defects in U. The sizes of v, v_1, v_2 and U then imply that C is contained in U and is safe for v in U.

The lemma then follows immediately from Corollary 6.6 (with U replacing A and the i-fringe replacing A'). □

8 Step D Proofs

Proof of Lemma 3.2. Consider Lemma 6.3 with A replaced by p'. Since $rl(v_1), rl(v_2) < \frac{m_1'}{8}$ and $cl(v_1), cl(v_2) < \frac{m_2'}{8}$, all points in p' except those in its top right corner sub-block of size $2\lfloor\frac{m_1'}{8}\rfloor \times 2\lfloor\frac{m_2'}{8}\rfloor$ and its bottom left corner of the same size are in $good_{p'}(v_1, v_2)$. p'' does not overlap either of these blocks. Therefore C does not overlap either of these blocks and all points in C are in $good_{p'}(v_1, v_2)$. The lemma then follows from Lemma 6.3. □

Proof of Lemma 3.3. By Lemma 3.2 and the fact that $rl(v_1), rl(v_2), rl(v) < \frac{m_1'}{8}$ and $cl(v_1), cl(v_2), cl(v) < \frac{m_2'}{8}$, C is safe for v in p'. The lemma then follows from Corollary 6.6 (with A replaced by p' and A' replaced by p). □

Proof of Lemma 5.1. By Step D.1a, any vector in S'' is clearly l-consistent and b-consistent in q. So assume that $v \in S'$.

First, suppose that the defect γ (see Fig.7) is not in the left $\lfloor\frac{m_2'}{8}\rfloor$ columns of q. It suffices to show that for all l-defects $a \in q$, $a + v$ is either outside q or an l-defect, and likewise for $a - v$. We only show that if $a + v \in q$ then $a + v$ is an l-defect; the other case is similar. Suppose for a contradiction that $a + v \in q$ is not an l-defect. Since v survives Step D.1b, $a + v$ must be a defect by Lemma 3.3. Let b be the l-defect in the row containing $a + v$. Since b is not in the left $\lfloor\frac{m'}{8}\rfloor$ columns of q and $b - v$ is in the row containing a, $b - v \in q$. Since $b - v$ is to the left of a, $b - v$ is not a defect. $(b - v, b)$ is thus a witness for v which would have been found in Step D.1b, which is a contradiction.

Next, suppose that γ in q is in the left $\lfloor\frac{m_2'}{8}\rfloor$ columns of q. To show that v is b-consistent in X, it suffices to show that for all b-defects $a \in X$, $a - v$ is either outside X or a b-defect and likewise for $a + v$. We only show that if $a - v \in X$ then $a - v$ is a b-defect; the other case is similar. Suppose for a contradiction that $a - v \in X$ is not a b-defect. Since v survives Step D.1b, $a - v$ must be a defect by Lemma 3.3. Let b be the b-defect in the column containing $a - v$. By Lemma 3.2, p'' has no defects. Therefore b is above p''. In addition, $b + v$ is in the column containing a, and thus $b + v \in X$. Since $b + v$ is below a, $b + v$ is not a defect. $(b, b + v)$ is thus a witness for v which would have been found in Step D.1b, which is a contradiction. □

Proof of Lemma 5.2. Let e be the top left corner of p''. Let C' be the cell bounded by $e, e + v, e + w, e + v + w$ in the (v, w)-lattice with respect to e. The sizes of v, w, p'' and Y imply that all points in C' are in p'', $good_X(v, w)$, and $good_Y(v, w)$. Lemma 3.2 implies that C' contains no defects with respect to C (C was defined in Section 3.4).

First, suppose that v, w are l-consistent in Y. By Lemma 6.4 (with Y replacing A and C, C' replacing C', C, respectively), all l-defects in Y are outside $good_Y(v, w)$. By Lemma 3.2, p'' has no defects. Since p'' has size $\frac{m_1'}{2} \times \frac{m_2'}{2}$ and $\frac{m_1'}{2} \geq rl(v) + rl(w) - 1$, $\frac{m_2'}{2} \geq cl(v) + cl(w) - 1$, the bottom left corner sub-block of Y of size $(rl(v) + rl(w) - 1) \times (cl(v) + cl(w) - 1)$ has no defects. Then, by the definition of $good_Y(v, w)$, all l-defects and hence all defects in Y are in Y'. The first part of the lemma follows.

Second, suppose for a contradiction that v, w are b-consistent in X and the leftmost defect α in q is in the left $\lfloor\frac{m_2'}{8}\rfloor$ columns of q. Note that $\alpha \in X$. By Lemma 3.2, all defects in X are above p''. Therefore, $\alpha \notin p''$ and $\alpha \in good_X(v, w)$. By Lemma 6.4 (with X replacing A and C, C' replacing C', C, respectively), all b-defects in X are outside $good_X(v, w)$. Since p'' has no defects, the definition of $good_X(v, w)$ implies that all b-defects and hence all defects in X

are in the top right corner sub-block of X of size $(rl(v)+rl(w)-1)\times(cl(v)+cl(w)-1)$. It follows that $\alpha \notin good_X(v,w)$, a contradiction. $\Box$

Proof of Lemma 5.4. Suppose for a contradiction that $w-v$ is not a Quad I vector. Since $rl(w)\geq rl(v)$, $v-w$ must be a Quad II vector. Recall that β is the leftmost bottommost defect in q. By Corollary 5.3, β and all other defects in q are in TR_q. Then, since $v-w$ is a Quad II vector, the sizes of v, w, TR_q imply that $\beta-v+w$ is in p and not a defect. Since S' is non-empty and $v\in S'$, recall from Case 2 of Step D.1a that $\beta-v$ is a l-defect in q. Since $\beta-v$ is a defect and $\beta-v+w$ is not a defect, $(\beta-v,\beta-v+w)$ is a witness for w by Lemma 3.3; further this witness would have been found in Step D.1b, contradiction. $\Box$

Proof of Lemma 3.4 That $v-u$ and $w-v$ are Quad I vectors follows from Lemma 5.4. Suppose for a contradiction that $v-u$ and $w-v$ are not parallel.

Let i be the number of rows in the jth row set. By the definitions of segments and row sets, $i > rl(w)-rl(u) = rl(w-v)+rl(v-u)$. Note that row br, the bottommost row in every row set, contains a defect. Also note that there are at least $rl(w)+1\geq rl(v)+1\geq rl(u)+1$ rows in q above the jth row set.

Let Y be the sub-block of q comprising all rows including and below the topmost row in the jth row set. Note that the topmost row in Y is one of the top br rows of q. Since all defects in q are in TR_q, $br\leq \frac{m_1}{8}$. The sizes of p'' and q imply that row br is completely above p''. It follows that Y has at least $\frac{m_1'}{2}$ rows.

Note that vectors $w-v$ and $v-u$ cannot both be l-consistent in Y; otherwise, by Lemma 5.2, all defects in Y are in the topmost $rl(w-v)+rl(v-u)-1$ rows of Y, which contradicts the fact that row br contains a defect.

It follows that for some l-defect $a\in Y$, one of $a+(w-v), a-(w-v), a+(v-u), a-(v-u)$ is in Y but not an l-defect. Assume that $a+(w-v)$ is in Y but not an l-defect; the other cases are similar. Since there are at least $rl(w)+1\geq rl(v)+1$ rows in q above Y and since a is in the rightmost $\lfloor\frac{m_2}{8}\rfloor$ columns of q, $a-v\in q$. Since v is l-consistent in q, $a-v$ is an l-defect. Since w is l-consistent in q, $a-v+w=a+(w-v)$ is an l-defect, a contradiction. $\Box$

This leads to the following theorem.

Theorem 8.1 *There exists a CRCW-PRAM algorithm which finds witnesses for all non-period vectors of an $m_1\times m_2$ pattern of row length less than $\lfloor\frac{m_1}{16}\rfloor$ and column length less than $\lfloor\frac{m_2}{16}\rfloor$; it takes $O(\log\log m)$ running time and does $O(m_1\times m_2)$ work.*

9 Computing witnesses for all Quad I and Quad II vectors

Given the algorithm for computing witnesses for all valid non-period Quad I and Quad II vectors, we show how to compute witnesses for all non-period Quad I and Quad II vectors. We describe the algorithm only for Quad I vectors; Quad II vectors are handled similarly.

Definitions. The *upper half* of a consecutive subset of i rows of p denotes the upper $\lfloor\frac{i}{2}\rfloor$ rows of this subset. The *left half* of a consecutive subset of columns of p is defined analogously.

We define $O(\log m_2)$ *c-sets* and $O(\log m_1)$ *r-sets* as follows. The first r-set comprises the upper half of the rows of p. The first c-set comprises the left half of the columns of p. r-set i, $1<i\leq\log m_1$ comprises the upper half of those rows of p which are below the $(i-1)$th r-set. Similarly, c-set i, $1<i\leq\log m_2$ comprises the left half of those columns of p which are to the

right of the $(i-1)$th c-set. Let r_i and c_i denote the number of rows of the ith r-set and the the number of columns of the ith c-set, respectively.

Sub-block (i,j) of p, $1 \leq i \leq \log m_1$, $1 \leq j \leq \log m_2$, is defined to be the $r_i \times c_j$ block formed by the intersection of the ith r-set and the jth c-set.

The *continuation* of a sub-block q of p is defined to be the sub-block of p with the same top left corner as q and the same bottom right corner as p. Note that the continuation of sub-block (i,j) has size at least $2r_i \times 2c_j$ and at most $(2r_i+1) \times (2c_j+1)$.

Let d_v denote the character at which the head of Quad I vector v lies when its tail is at the top left corner of p. We redefine the term *fall* as follows in this section. v is said to *fall* at d_v; in addition if d_v is in sub-block q of p, v is said to *fall* in q. Note that if vector v falls in sub-block q of p, any witness for v in p must have its head in the continuation of q.

The Block Structure. All sub-blocks (i,j), $1 \leq i \leq \log m_1$, $1 \leq j \leq \log m_2$, are processed in parallel. Consider sub-block (i,j) of p. Let q denote this sub-block. q is further partitioned into a constant number of sub-blocks of size $\lfloor \frac{r_i}{16} \rfloor \times \lfloor \frac{c_j}{16} \rfloor$ (except at the boundary of q where truncated sub-blocks are used). All such sub-blocks are processed in parallel. Let q' be one such sub-block. For all Quad I vectors of p which fall in q', a witness with head at a character in the continuation of q', if any, is found using the procedure described in the rest of this section. The processing of q' takes $O(\log \log m)$ time and $O(|q|)$ work. Thus the work done over all sub-blocks of q is $O(|q|)$ and the work done over all sub-blocks (i,j) of p is $O(|p|)$. The total time taken is $O(\log \log m)$.

We need the following definitions before describing the processing of sub-block q' of q.

Definitions. The *(i,j)-prefix* of p is the sub-block of p of size $r_i \times c_j$ with the same top left corner as p. Let U be the (i,j)-prefix of p, E the continuation of q', and L the maximal sub-block of E completely below and to the right of q'. See Fig.14. Note that $|E| = O(|q|)$. Let V be the set of Quad I vectors of p which fall in q'.

The *region* of a vector v is defined to be the sub-block of E whose top left corner is d_v and whose bottom right corner is the bottom right corner of E. A witness $(a,b) \in E$ for a Quad I or Quad II vector $w = v - u$, $u, v \in V$, is said to be *good* for u and v if a is in the region of u (and therefore, b is in the region of v). See Fig.13. Note that v and u can be duelled using the witness (a,b) if and only if witness (a,b) is good for u and v.

There are two main steps.

Step 1. Witnesses are found for some of the vectors in V in this step in $O(\log \log m)$ time and $O(|E|)$ work. Two vectors $u, v \in V$ survive this step only if there is no witness in E which is good for u and v. Clearly, if u and v survive this step then the two copies of p placed with top left corners at d_u and d_v match wherever both overlap E.

Step 2. Witnesses, if any, are found for the surviving vectors in V in $O(1)$ time and $O(|E|)$ work. This is done using the procedure $Verify(E, E', W)$ (see Appendix), where E' is the sub-block of the pattern with the same size as E and the same top left corner as p, and W is the set $\{d_v | v \text{ is a surviving vector in } V\}$. By Step 1, copies of p and hence copies of E' placed at the locations in V match each other wherever they overlap; the input conditions to $Verify(E, E', W)$ are therefore satisfied. If bad_{d_v} exists for some $d_v \in V$, then $(bad_{d_v} - v, bad_{d_v})$ is a witness for v.

Next, we describe Step 1 in detail. There are 4 substeps.

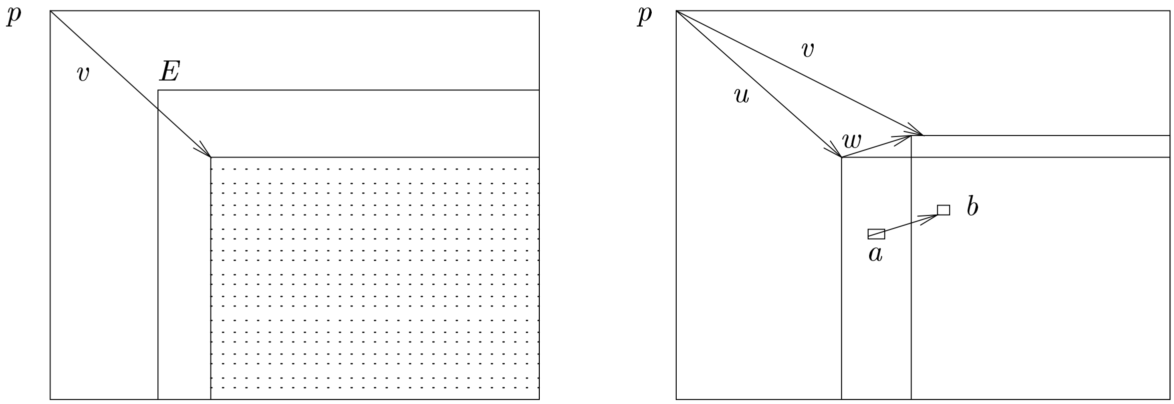

a) The region of vector v b) (a, b) is good for u and v

Figure 13:

Step 1.1. Witnesses in U are found for all non-period Quad I and Quad II vectors of U of row length less than $\lfloor \frac{r_i}{16} \rfloor$ and column length less than $\lfloor \frac{c_j}{16} \rfloor$. By Theorem 8.1, the time taken for this is $O(\log \log m)$ and the work done is $O(|U|) = O(|E|)$. Similarly, witnesses in L are found for all non-period Quad I and Quad II vectors of L with row length less than $\lfloor \frac{r_i}{16} \rfloor$ and column length less than $\lfloor \frac{c_j}{16} \rfloor$ in $O(\log \log m)$ time and $O(|L|) = O(|E|)$ work.

Step 1.2. Witnesses are found for some of the surviving vectors in V in this step. Vector $v \in V$ survives this step only if there is an occurrence of U with top left corner at d_v; otherwise, any point a at which this copy of U mismatches E is the head of a witness for v. This step is accomplished in $O(\log \log m)$ time and $O(|E|)$ work as follows.

Note that since U has size $r_i \times c_j$, q has size $r_i \times c_j$ and E has size at least $2r_i \times 2c_j$, any copy of U with top left corner in q' lies completely within E. Recall that witnesses in U, if any, were computed in Step 1 for all Quad I and Quad II vectors of row length less than $\lfloor \frac{r_i}{16} \rfloor$ and column length less than $\lfloor \frac{c_j}{16} \rfloor$. Using these witnesses, all occurrences of U with top left corner in q' are found using the pattern matching algorithm of [ABF93]. In addition, for each vector $v \in V$, if a copy of U placed with top left corner at d_v mismatches E at some location, one such point of mismatch is computed. This point is the head of a witness for v.

Step 1.3. Witnesses are found for some of the surviving vectors in V in this step. Let v_0 be the vector which falls at the bottom right corner of q'. Vector $v \in V$ survives this step only if there is an occurrence of L with top left corner at $d_{v_0 - v}$. This step is accomplished in $O(\log \log m)$ time and $O(|E|)$ work as follows.

All occurrences of L with top left corner in U are found in $O(\log \log m)$ time and $O(|E|)$ work using the pattern matching algorithm of [ABF93]. The witnesses computed in Step 1 are used in this process. If the copy of L with top left corner at $d_{v_0 - v}$ mismatches the portion of p it overlaps at some location d, then $(d, d + v)$ is a witness for v.

Step 1.4. There are three cases, depending upon whether or not all the difference vectors of the surviving vectors in V are parallel. All difference vectors are parallel Quad I vectors

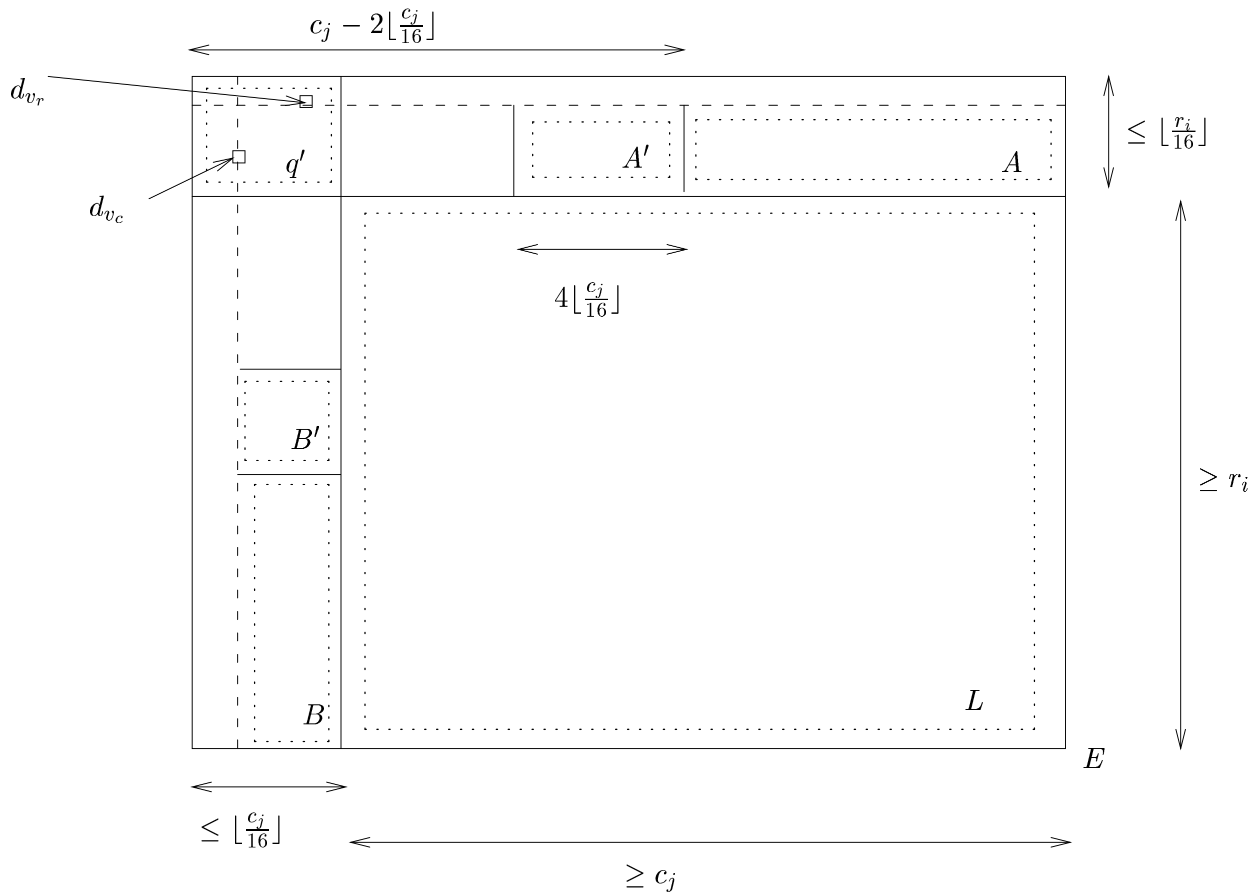

Figure 14: The sub-blocks A, A', B, B', L.

in Case 1 and parallel non-horizontal, non-vertical Quad II vectors in Case 2. In the third case, there are at least two non-parallel difference vectors. In all three cases, vectors $u, v \in V$ survive this step only if there is no witness in E which is good for u and v.

Which of these three cases holds can be determined in $O(\log\log m)$ time and $O(|q'|) = O(|E|)$ work as follows. For each surviving vector $v \in V$, the difference vector of v and a particular surviving vector $w \in V$ is computed in $O(1)$ time and $O(|q'|) = O(|E|)$ work. The slope of all these difference vectors are compared with the slope of any arbitrary one of these difference vectors, v_1 say. If all these difference vectors are parallel to v_1, then the slope of v_1 determines which of Case 1 and Case 2 hold. Otherwise, if one of these difference vectors v_2 is not parallel to v_1 then Case 3 holds. Each case is processed in $O(\log\log m)$ time and $O(|E|)$ work.

The following definitions are helpful.

Definitions. Let V' be the set of vectors in V which survive Step 1.3. Without loss of generality, assume that $|V'| \geq 2$. Let V_{dif} be the set of all Quad I and Quad II difference vectors of the vectors in V'. For each $v \in V'$, let U_v denote the copy of U placed with top left corner at d_v.

See Fig.14. Let v_r and v_c be the vectors in V' with minimum row and column lengths, respectively. Let A denote the sub-block of E which is bounded below by L, above by (and including) the row containing d_{v_r}, to the left by the column $c_j - 2\lfloor \frac{c_j}{16} \rfloor$ of E, and to the right by the right boundary of E. Let B denote the sub-block of E which is bounded to the right by L, to the left by (and including) the column containing d_{v_c}, above by the row $r_i - 2\lfloor \frac{r_i}{16} \rfloor$ of E, and below by the bottom boundary of E. The following fact is easily seen from the definitions of A, B and U.

Fact 1. For any $v \in V'$, U_v has at least $2\lfloor \frac{c_j}{16} \rfloor$ columns to the right of the left boundary of A and at least $2\lfloor \frac{r_i}{16} \rfloor$ rows below the upper boundary of B.

Let A' be the sub-block to the left of A with the same rows as A and $4\lfloor \frac{c_j}{16} \rfloor$ columns. Let B' be the sub-block above B with the same columns as B and $4\lfloor \frac{r_i}{16} \rfloor$ rows.

Lemma 9.1 *For all $w \in V_{dif}$, w does not have a witness in either U or L.*

Proof. Suppose $w = v - u$, where $u, v \in V'$. By Step 1.2, there is an occurrence of U with top left corner at d_u and another with top left corner at d_v. It follows that w is a period vector of U. By Step 1.3, there is an occurrence of L with top left corner at $d_{v_0 - u}$ and another with top left corner at $d_{v_0 - v}$, where v_0 is the vector which falls at the bottom right corner of q'. It follows that w is a period vector of L. □

Lemma 9.2 implies that only those witnesses for vectors in V_{dif} which have at least one endpoint in one of A or B are of interest.

Lemma 9.2 *Let $u, v \in V'$, and let $w = v - u$. Any witness in E for w which is good for u and v has at least one endpoint in either A or B. In particular, if w is a Quad I vector then the tail of such a witness is either in A or in B, and if w is a Quad II vector, either the tail of such a witness is in B or the head is in A.*

Proof. Let (a, b) be a witness for w in E which is good for u and v.

First, we show that one of a, b is either in A or in B. Note that since w is a Quad I or Quad II vector, b must be either vertically aligned with or to the right of a. Further, since

(a, b) is good for u and v, a must be in the region of u. The region of u is covered by four overlapping sub-blocks, U_u, A, B, L. If $a \in U_u$ then $b \in U_v$; thus a cannot be within U_u, as otherwise at most one of U_u, U_v would completely match E, which contradicts the outcome of Step 1.2. If $a \notin L \cup U_u$ then a is in either A or B. The only remaining case is when $a \in L - U_u$. By Lemma 9.1, a and b cannot both be in L. Since b is aligned with or to the right of a and $b \notin L$, Fact 1 implies that b is in A in this case.

From the above paragraph, when w is a Quad I vector, a is in A or in B except in the case when $a \in L - U_u$. In this case, $b \in L$, which contradicts Lemma 9.1; so this case does not arise for Quad I vectors.

Finally, consider the case when w is a Quad II vector. Since a is in the region of u, b is in the region of v; therefore, neither a nor b can be above A or to the left of B. It follows that if $b \in B$ then $a \in B$, and if $a \in A$ then $b \in A$. Since one of a, b is in A or in B, either $a \in B$ or $b \in A$. □

Case 1. All vectors in V_{dif} are parallel and are Quad I vectors.

$Line(V', E, p)$ completes the algorithm in this case (see Section 2 for the definition of $Line$). This takes $O(\log \log m)$ time and $O(|E|)$ work, as claimed.

Case 2. All vectors in V_{dif} are parallel and are non-vertical, non-horizontal Quad II vectors.

Procedure $Line$ cannot be used here as it requires that the vectors involved and their difference vectors both be either Quad I vectors or Quad II vectors, which is not true in this case.

Let w_{min} be the difference vector with the smallest row length. The bottommost point $e \in A$, if any, such that $e \not\equiv e - w_{min}$ and the rightmost point $f \in B$, if any, such that $f \not\equiv f + w_{min}$ are found in $O(1)$ time and $O(|E|)$ work. Lemma 9.3 shows that for all $w \in V_{dif}$, $(e - w, e)$ and $(f, f + w)$ are witnesses for w.

Next, all pairs of vectors in V' are duelled in parallel using the witnesses computed above. Since at most one vector in V' falls in each row and column of q', $(|V'|)^2 \leq |E|$, and therefore, the total work done is $O(|E|)$ and the time taken is $O(1)$.

We claim that if one of e or f exists and is in the region of either u or v, $u, v \in V'$, $rl(v) \leq rl(u)$, then one of u, v is eliminated when the two are duelled. To see this, suppose that e exists and is in the region of either u or v (the case when f exists is similar). Let $w = v - u$; note that w is a Quad II vector in V_{dif}. Since $e \in A$, e is in the region of v and $e - w$ is in the region of u. Therefore, the witness $(e - w, e)$ is good for u and v (see Fig.13) and one of u, v is eliminated by duelling using this witness. So vectors $u, v \in V'$ survive this step only if their regions are below e, if e exists, and to the right of f, if f exists. Then, by Lemma 9.4, there is no witness for $v - u$ in E which is good for u and v. This completes Step 1.4 for this case. This case requires $O(1)$ time and $O(|E|)$ work, as claimed.

Lemma 9.3 *For all $w \in V_{dif}$, $e \not\equiv e - w$ and $f \not\equiv f + w$.*

Proof. Recall that $rl(w_{min}), rl(w) < \lfloor \frac{r_i}{16} \rfloor$, $cl(w_{min}), cl(w) < \lfloor \frac{c_j}{16} \rfloor$, q' has size at most $\lfloor \frac{r_i}{16} \rfloor \times \lfloor \frac{c_j}{16} \rfloor$ and U has size $r_i \times c_j$. Further, recall that $|V'|$ is non-empty and for all $v \in V'$, U_v matches E.

We show that $e \not\equiv e - w$. $f \not\equiv f + w$ is shown similarly. By the definition of A and the sizes of w_{min}, w and U, there exists an integer k such that $e - kw_{min}$ and $e - kw_{min} - w$ are both either in L or U_{v_r}. By Lemma 9.1, $e - kw_{min} \equiv e - kw_{min} - w$. Since w_{min} is a non-horizontal Quad II vector, the definition of e, Lemma 9.1, and Fact 1 together imply that $e \not\equiv e - w_{min} \equiv e - kw_{min} \equiv e - kw_{min} - w$ and $e - w \equiv e - w - kw_{min}$. The claim follows. □

Lemma 9.4 *Let u, v be vectors in V' and $w = v - u \in V_{dif}$. Suppose that either e does not exist or the regions of u, v are both below e and that either f does not exist or the regions of u, v are both to the right of f. There is no witness for w in E which is good for u and v.*

Proof. We show that for all a in the region of u such that $a + w \in A$, $a \equiv a + w$. A similar argument shows that for all a in the region of u such that $a \in B$, $a \equiv a + w$. Since w is a Quad II vector, it follows that there is no witness for w which is good for u and v with either tail in B or head in A. By Lemma 9.2, there is no witness for w in E which is good for u and v; the lemma follows.

Consider any a in the region of u such that $b = a + w$ is in A. Clearly, b is in the region of v. Since e, whenever it exists, is above the region of v, it is above b. As in the the proof of Lemma 9.3, there is an integer k such that $b - kw_{min}, b - kw_{min} - w$ are both either in L or U_v and $b - kw_{min} \equiv b - kw_{min} - w$. Since e either does not exist or is above b, the definition of e, Lemma 9.1, and Fact 1 together imply that $b \equiv b - kw_{min} \equiv b - kw_{min} - w$ and $b - w \equiv b - w - kw_{min}$. Therefore, $a \equiv b = a + w$, as claimed. □

Case 3. There are at least two non-parallel difference vectors, v_1 and v_2, in V_{dif}.

We show how to find witnesses for some of the surviving vectors in V' in $O(\log \log m)$ time and $O(|E|)$ work with the result that if $u, v \in V'$ survive then there is no witness with an endpoint in A which is good for u and v. A similar elimination procedure can be performed to ensure that if $u, v \in V'$ survive then there is no witness with an endpoint in B which is good for u and v. Recall that, by Lemma 9.2, any witness which is good for u and v, $u, v \in V'$, has one endpoint in A or B. It follows that if $u, v \in V'$ survive then there is no witness in E which is good for u, v. The goal of Step 1 is thus accomplished in $O(\log \log m)$ time and $O(|E|)$ work, as claimed.

Definitions. Let g be the point in row $8\lceil \frac{r_i}{16} \rceil$ and column $8\lceil \frac{c_j}{16} \rceil$ of E. Let C be the cell in the (v_1, v_2)-lattice with respect to point g and bounded by the points $g, g + v_1, g + v_2, g + v_1 + v_2$. Recall that $rl(v_1), rl(v_2) < \lfloor \frac{r_i}{16} \rfloor$, $cl(v_1), cl(v_2) < \lfloor \frac{c_j}{16} \rfloor$, U has size $r_i \times c_j$, and L has size at least $r_i \times c_j$.

For a sub-block S of E, let BR_S, BL_S, TR_S and TL_S denote the bottom right, bottom left, top right and top left sub-blocks, respectively, of S, of size $2\lfloor \frac{r_i}{16} \rfloor \times 2\lfloor \frac{c_j}{16} \rfloor$. These sub-blocks of S are called *corner blocks* of S.

The following lemmas are useful.

Lemma 9.5 *All defects in L with respect to C are in BR_L, BL_L, TR_L, and TL_L. In addition, there are no defects with respect to C in A' or B'.*

Proof. The sizes of v_1, v_2, L, U and the definition of C imply that all points in C are in $good_L(v_1, v_2)$, $good_{U_{v_r}}(v_1, v_2)$ and $good_{U_{v_c}}(v_1, v_2)$. By Lemma 6.3 applied to L and U (i.e., U_{v_r} or U_{v_c}), all defects in L with respect to C are in BR_L, BL_L, TR_L and TL_L, and all defects in U with respect to C are in BR_U, BL_U, TR_U, and TL_U.

The sizes of v_1, v_2 along with Fact 1 imply that none of the corner blocks of U_{v_r} or U_{v_c} intersect with A' or B'. Then, since A' is contained in U_{v_r}, there are no defects with respect to C in A'. Since B' is contained in U_{v_c}, there are no defects with respect to C in B'. □

Lemma 9.6 *(a, b) is a witness in E for a vector $w \in V_{dif}$ if and only if either exactly one of a, b is a defect with respect to C, or both a, b are defects with respect to C and the characters at a and b are different.*

Proof. The sizes of w, v_1, v_2, L imply that for all points $e' \in C$, $e' + w \in good_L(v_1, v_2)$. Therefore C is safe for w in L. Further, by Lemma 9.1, w is a period vector of L. The lemma follows from Corollary 6.6, with L and E replacing A and A', respectively. □

Definition. Let α be the leftmost bottommost defect with respect to C in $A \cup TR_L$. Let G be the rectangle bounded above by the upper boundary of A, below by the row containing α, to the left by the left boundary of A' and to the right by the right boundary of E.

Lemma 9.7 *The only defects with respect to C in G are in $G \cap (A \cup TR_L)$. In particular, there are no defects with respect to C in the left $4\lfloor \frac{c_j}{16} \rfloor$ columns of G.*

Proof. Both $G - (A \cup TR_L)$ and the left $4\lfloor \frac{c_j}{16} \rfloor$ columns of G comprise characters in A' and $L - TR_L - TL_L - BL_L - BR_L$; by Lemma 9.5, none of these characters is a defect with respect to C. □

The following 5 substeps are performed in Step 1.4 in this case. Each takes $O(\log\log m)$ time and $O(|E|)$ work, as claimed.

Step 1.4a. α and all l-defects and b-defects in G are found in $O(1)$ time and $O(|E|)$ work. Lemmas 9.8 and 9.9 hold following this step.

Lemma 9.8 *For all Quad II vectors $w \in V_{dif}$, $\alpha - w$ is in E and $\alpha - w \not\equiv \alpha$.*

Proof. The length of w implies that $\alpha - w$ is either in A, A' or L. First, suppose $\alpha - w \in A \cup A'$. By Lemma 9.5, A' has no defects with respect to C; the lemma follows from the definition of α and from Lemma 9.6. Second, suppose $\alpha - w \in L$. Since $\alpha \in A \cup TR_L$, the length of w implies that $\alpha - w \in L - TL_L - BL_L - BR_L$. By the definition of α and the fact that w is a Quad II vector, $\alpha - w$ cannot be a defect with respect to C in TR_L. Then, by Lemma 9.5, $\alpha - w$ is not a defect. The lemma follows from Lemma 9.6. □

Lemma 9.9 *Let $u, v \in V'$. Suppose either α exists and is above d_u and d_v or α does not exist. Then there is no witness in E with an endpoint in A which is good for u and v.*

Proof. Suppose for a contradiction that there a witness (a, b) with one endpoint in A which is good for u and v. Clearly, both a and b are below α, if it exists, and in either A, A' or $L - TL_L - BR_L - BL_L$. By the definition of α, neither a nor b is a defect with respect to C in TR_L. Then, by Lemma 9.5, neither a nor b is a defect with respect to C. Lemma 9.6 then gives the required contradiction. □

Recall that we are looking for witnesses with at least one endpoint in A (see beginning of Case 3). There are two cases next, depending upon the existence of α. First, suppose α does not exist. By Lemma 9.9, for no $u, v \in V'$ is there a witness which has an endpoint in A and is good for u, v. Nothing further needs to be done in this case. Second, suppose α exists. The set of surviving vectors in V' is partitioned into two classes. The first class contains those vectors v such that α is aligned with or below d_v. The second class contains those vectors v such that α is above d_v. In subsequent steps, we show how witnesses in L, if any, are found for all vectors in the first class. Thus only vectors in the second class survive these steps. By Lemma 9.9, for vectors u, v in the second class, there is no witness good for u, v with an endpoint in A; this accomplishes our goal.

Definition. The rest of Step 1.4 assumes that α exists. V' is now redefined to contain only those surviving vectors $v \in V$ such that α is aligned with or below d_v. V_{dif} is also redefined to be the set of all Quad I and Quad II difference vectors of the vectors in V'. By Lemma 9.8, V' has the property that if $u, v \in V'$ and $v - u$ is a Quad II vector then $(\alpha - v + u, \alpha)$ is a good witness for u and v.

Step 1.4b. Witnesses will be found for all but at most $\min\{\lfloor \frac{r_i}{16} \rfloor, \lfloor \frac{c_j}{16} \rfloor\}$ of the vectors in V' in this step. Further, if $u, v \in V'$ survive where $rl(v) \geq rl(u)$, $v - u$ is a non-horizontal non-vertical Quad I vector. The total work done in this step is $O(|q'|) = O(|E|)$ and the total time taken is $O(\log \log m)$.

First, all rows of q' are processed in parallel. Consider row r. All $v \in V'$ which fall in r are considered. In $O(\log \log m)$ time and $O(|r|)$ work, witnesses are found for all but one of these vectors by duelling amongst them. This is done using a procedure similar to the one used in Substep B.1. Only one vector which falls in r survives because for any two vectors $u, v \in V'$ which fall in r with $cl(v) > cl(u)$, $v - u$ is horizontal and hence a Quad II vector and by Lemma 9.8, a witness for $v - u$ which is good for u and v was found in Step 1.4a.

Second, all columns of q' are processed in parallel in a similar manner. Consider column c. All $v \in V'$ which fall in c are considered. In $O(\log \log m)$ time and $O(|c|)$ work, witnesses are found for all but one of these vectors by duelling amongst them.

Finally, all pairs of surviving vectors in V' are duelled. Since at most one vector falls in each row and column of q', the number of pairs and hence the work done is $O(|q'|) = O(|E|)$. Recall from Step 1.4a that for any particular pair u, v, if $v - u$ is a Quad II vector, then a witness for $v - u$ which is good for u and v is known. u and v are duelled using this witness in $O(1)$ time.

Step 1.4c. Witnesses are found for some of the surviving vectors in V' in this step. Following this step, there will be at most $\min\{\lfloor \frac{r_i}{16} \rfloor, \lfloor \frac{c_j}{16} \rfloor\}$ distinct difference vectors for the surviving vectors in V'. This is accomplished in $O(1)$ time and $O(|q'|) = O(|E|)$ work in two steps.

First, all difference vectors $w \in V_{dif}$ such that $v - u = w$ are processed in parallel, where u, v are surviving vectors in V' and $rl(v) > rl(u)$. By Step 1.4b, these difference vectors are non-horizontal non-vertical Quad I vectors. Since at most one surviving vector in V' falls in each row and each column of q', there are at most $O(|q'|)$ such vectors w. Witnesses are found for some of these vectors in this step in $O(1)$ time and $O(|q'|)$ work using the procedure implied by Lemma 9.10. Further, as shown in Lemma 9.10, the witness found for vector $w = v - u$ is good for u and v.

Second, all pairs of surviving vectors in V' are duelled using the witnesses computed above in $O(1)$ time and $O(|q'|)$ work. By Lemma 9.10, if $u, v \in V'$, $rl(v) > rl(u)$, survive then $\alpha - (v - u)$ is both an l-defect and a b-defect in G; in addition, since $v - u$ is a non-horizontal non-vertical Quad I vector by Step 1.4b, $cl(v) > cl(u)$. It follows that among the difference vectors of the surviving vectors in V', there is at most one vector for each row length less than $\lfloor \frac{r_i}{16} \rfloor$ and at most one vector for each column length less than $\lfloor \frac{c_j}{16} \rfloor$. Therefore, there are at most $\min\{\lfloor \frac{r_i}{16} \rfloor, \lfloor \frac{c_j}{16} \rfloor\}$ distinct difference vectors of the surviving vectors in V'.

Lemma 9.10 *Suppose that u and v are surviving vectors in V' with $rl(v) > rl(u)$. Let $w = v - u$. $\alpha - w$ is in the region of u and in G. Further, if $\alpha - w$ is not both an l-defect and a b-defect in G then either $(\alpha, \alpha - w)$, $(\alpha_l, \alpha_l + w)$, or $(\alpha_b, \alpha_b + w)$ is a witness for w, where α_l is the l-defect in G in the row containing $\alpha - w$ and α_b is the b-defect in G in the column containing $\alpha - w$.*

Proof. Recall that $\alpha \in A \cup TR_L$ and that, by Step 1.4a, d_u, d_v are aligned with or above α. Clearly, d_u is aligned with or above $\alpha - w$. Moreover, since $cl(w) < \frac{c_j}{16}$, $\alpha - w \in G$ and therefore, d_u is to the left of $\alpha - w$. It follows that $\alpha - w$ is in the region of u.

If $\alpha - w$ is not a defect, the lemma follows from Lemma 9.6. Suppose $\alpha - w$ is a defect but not both an l-defect and a b-defect in G. We consider only the case when $\alpha - w$ is not an l-defect in G and show that $\alpha_l + w$ is not a defect in this case. In the case when $\alpha - w$ is not a b-defect, a similar argument shows that $\alpha_b + w$ is not a defect. In both cases, the lemma follows from Lemma 9.6.

Suppose for a contradiction that $\alpha_l + w$ is a defect. As shown earlier, $\alpha - w \in G$. Since $\alpha - w$ is a defect, by Lemma 9.7, $\alpha - w \in A \cup TR_L$ and therefore $\alpha_l \in A \cup TR_L$. Since w is a Quad I vector, the length of w implies that $\alpha_l + w \in A \cup L - TL_L - BR_L - BL_L$. Since $\alpha_l + w$ is a defect, $\alpha_l + w \in A \cup TR_L$ by Lemma 9.5. Further α_l is to the left of $\alpha - w$ and therefore $\alpha_l + w$ is to the left of α. This contradicts the fact that α is the bottommost leftmost defect in $A \cup TR_L$. □

Definition. For a vector $v \in V'$, let G_v be the portion of G aligned with or below d_v.

Let $\hat{V}_{dif}$ to be the set of vectors w such that $w = v - u$ for some surviving vectors $u, v \in V'$ with $rl(v) > rl(u)$. Note that $|\hat{V}_{dif}| \leq \min\{\lfloor \frac{r_i}{16} \rfloor, \lfloor \frac{c_j}{16} \rfloor\}$ and that all vectors in $\hat{V}_{dif}$ are non-horizontal, non-vertical Quad I vectors by Step 1.4b.

A Quad I vector w is said to be *strongly l-consistent* in a sub-block X of E if for all l-defects $\gamma \in X$, $\gamma + w$ is either an l-defect in X, a defect in E outside X, or outside E and $\gamma - w$ is either an l-defect in X or outside X.

Step 1.4d. Witnesses are found for some of the surviving vectors in V' is this step. If vectors $u, v \in V'$, $rl(v) > rl(u)$, survive this step then the vector $w = v - u$ is strongly l-consistent in G_u. This is done in two stages.

First, all vectors in $\hat{V}_{dif}$ are processed in parallel. Consider vector $w \in \hat{V}_{dif}$. Of all witnesses for w, if any, with one endpoint being a l-defect in G, the witness with the bottommost head is found (with ties broken arbitrarily). Let $(\beta_w, \beta_w + w)$ denote this witness. This takes $O(1)$ time and $O(r_i)$ work per vector in $\hat{V}_{dif}$. Since $|\hat{V}_{dif}| \leq \min\{\lfloor \frac{r_i}{16} \rfloor, \lfloor \frac{c_j}{16} \rfloor\}$, the total work done is $O(|E|)$.

Second, all pairs of vectors $u, v \in V'$ with $rl(v) > rl(u)$ are processed in parallel in $O(1)$ time and $O(|q'|) = O(|E|)$ work. If a good witness for u and v was computed in the first stage then u and v are duelled using this witness.

Lemma 9.11 holds after this stage. Step 1.4d takes $O(1)$ time and does $O(|E|)$ work, as claimed.

Lemma 9.11 *If $u, v \in V'$, $rl(v) > rl(u)$, survive Step 1.4d then $w = v - u$ is strongly l-consistent in G_u.*

Proof. Recall from Step 1.4b that w is a Quad I vector. Suppose for a contradiction that w is not strongly l-consistent in G_u, i.e., there exists an l-defect γ in G_u such that either $\gamma - w \in G_u$ but is not an l-defect or $\gamma + w$ is in G_u but not an l-defect or $\gamma + w$ is in $E - G_u$ but not a defect. We consider each case in turn. In each case, we show that for some l-defect $\delta \in G_u$, δ is one endpoint of a witness (a, b) for w which is good for u, v. It follows that β_w exists (see first stage of Step 1.4d). and is aligned with or below a. Since $\beta_w \in G$, β_w is to the right of q'. It follows that β_w is in the region of u and therefore, $(\beta_w, \beta_w + w)$ is good for u and v.

Therefore, u and v would have been duelled using this witness in the second stage of Step 1.4d and one of u, v would have been eliminated, a contradiction.

First, suppose $\gamma - w \in G_u$ but is not an l-defect. If $\gamma - w$ is a non-defect then by Lemma 9.6, $(\gamma - w, \gamma)$ is a witness for w which is good for u and v. If $\gamma - w$ is a defect, then let δ be the l-defect in G_u in the row containing $\gamma - w$. Clearly, $\delta + w$ is a non-defect in G_u and by Lemma 9.6, $(\delta, \delta + w)$ is a witness for w which is good for u and v.

Second, suppose $\gamma + w \in G_u$ is not an l-defect. If $\gamma + w$ is a non-defect then by Lemma 9.6, $(\gamma, \gamma + w)$ is a witness for w which is good for u and v. If $\gamma + w$ is a defect then let δ be the l-defect in G_u in the row containing $\gamma + w$. Since, by Lemma 9.7, there are no defects with respect to C in the left $4\lfloor \frac{c_j}{16} \rfloor$ columns of G, δ is to the right of these columns and therefore, $\delta - w \in G_u$. Clearly, $\delta - w$ is a non-defect and by Lemma 9.6, $(\delta - w, \delta)$ is a witness which is good for u and v.

Third, suppose $\gamma + w$ is a non-defect in $E - G_u$. By Lemma 9.6, $(\gamma, \gamma + w)$ is a witness which is good for u and v. $\Box$

Step 1.4e. Witnesses with heads in G, if any, are found for all surviving vectors in V' in $O(\log \log m)$ time and $O(|E|)$ work. Lemma 9.13 shows that if $u, v \in V'$ survive this step then there is no witness with an endpoint in A which is good for u and v.

Step 1.4e is performed as follows. Consider the row partition of G (see Section 5 for definition of row partition). A vector $v \in V'$ is said to *fall* in segment k of G if the row of E containing d_v passes through segment k in G. All $O(\log r_i)$ segments of G are processed in parallel. Consider the kth segment. Let V'' be the set of surviving vectors $v \in V'$ which fall in this segment. By Lemma 9.12, all difference vectors of vectors in V'' are parallel. Note that any witness (a, b) for $v \in V''$ has the property that if $b \in G$ then $b \in K \cap G$, where K comprises those rows of E whose portions in G comprise the kth segment and the kth row set in G. $Line(V'', K, p)$ is used to find witness with heads in K, if any, for vectors in V'' (see Section 2 for the definition of *Line*). This takes $O(\log \log m)$ time and $O(|K|)$ work. Thus, the work done over all segments is $O(|E|)$. This completes Step 1.4e.

Lemma 9.12 *If surviving vectors $u, v, w \in V'$ fall in the kth segment, $rl(u) < rl(v) < rl(w)$, then $v - u$ and $w - v$ are parallel.*

Proof. Recall that $rl(v-u) + rl(w-v) < \lfloor \frac{r_i}{16} \rfloor$ and $cl(v-u) + cl(w-v) < \lfloor \frac{c_j}{16} \rfloor$.

By Lemma 9.11, $v - u$ and $w - v$ are strongly l-consistent in G_u and G_v, respectively. Let R be the sub-block comprising the rows of the kth row set. Since $R \subset G_u, G_v$, both $v - u$ and $w - v$ are strongly l-consistent in R. Suppose for a contradiction that $v - u$ and $w - v$ are non-parallel.

By Step 1.4b, $v-u$ and $w-v$ are both Quad I vectors. Note that R has at least $rl(v-u) + rl(w-v) + 1$ rows and at least $6\lfloor \frac{c_j}{16} \rfloor \geq 6(cl(v-u) + cl(w-v) + 1)$ columns. Let e be the topmost point in the $2\lfloor \frac{c_j}{16} \rfloor$th column from left in R. Let C' be the cell in the $(v-u, w-v)$-lattice with respect to e bounded by the points $e, e+v-u, e+w-v, e+w-u$. Since $v-u$ and $w-v$ are Quad I vectors, the sizes of $v-u$ and $w-v$ imply that C' is completely contained in the right half of the left $4\lfloor \frac{c_j}{16} \rfloor$ columns of R. By Lemma 9.7, there are no defects with respect to C in C'. In addition, all points in C' are outside the right and left $\lfloor \frac{c_j}{16} \rfloor > cl(v-u) + cl(w-v)$ columns of R. Therefore, all points in C' are in $good_R(v-u, w-v)$. By Lemma 6.4 (with $R, E, v-u, w-v, C, C'$ replacing A, A', v, w, C', C, respectively), all l-defects in R are outside $good_R(v-u, w-v)$. Since α is in the bottommost row of R and, by Lemma 9.7, to the right of the left $4\lfloor \frac{c_j}{16} \rfloor$ columns of R, α is in $good_R(v-u, w-v)$. This is a contradiction. $\Box$

Lemma 9.13 *Let u, v be vectors in V' which survive Step 1.4d, with $w = v - u$, $rl(v) > rl(u)$. Suppose there is a witness (a, b) for w which is good for u and v and has an endpoint in A. Either u or v has a witness with head in G.*

Proof. It suffices to show that both a and b are in G. Recall that $rl(w) < \lfloor \frac{r_i}{16} \rfloor$ and $cl(w) < \lfloor \frac{c_j}{16} \rfloor$. In addition, by Lemma 9.6, at least one of a, b must be a defect.

First, suppose for a contradiction that $a \in G$ and $b \notin G$. Since w is a Quad I vector, b must be below G. The size of w implies that $b \in A \cup A' \cup (L - TL_L - BR_L - BL_L)$. By Lemma 9.5 and the definition of α, b is not a defect. Then a must be a defect. Let h be the l-defect in G in the row containing a. As in the case of b, $h + w$ can be shown to be a non-defect below G. It follows that w is not strongly l-consistent in G, which contradicts Lemma 9.11.

Second, suppose for a contradiction that $b \in G$ and $a \notin G$. Since (a, b) is good for u and v, a is below d_u and hence below d_{v_r}. Since w is a Quad I vector, a must be to the left of G. The size of w implies that b is in the left $\lfloor \frac{c_j}{16} \rfloor$ columns of G. Then $(a, b) \in U_{v_r}$ which contradicts Lemma 9.1.

Finally, suppose for a contradiction that neither a nor b is in G. Since either a or b is in A and w is a Quad I vector, the size of w implies that $a, b \in A \cup A' \cup (L - TL_L - BR_L - BL_L)$. By Lemma 9.5 and the definition of α, neither a nor b is a defect, which contradicts Lemma 9.6. □

We conclude with the following theorem.

Theorem 9.14 *There exists a CRCW-PRAM algorithm which finds witnesses for all non-period Quad I and Quad II vectors of an $m_1 \times m_2$ pattern in $O(\log \log m)$ running time and does $O(m_1 \times m_2)$ work, where $m = \max\{m_1, m_2\}$.*

References

[AB92] A. Amir, G. Benson. Two dimensional periodicity in rectangular arrays. *Proc. of the 3rd ACM-SIAM Symposium on Discrete Algorithms*, 1992, pp. 440–452.

[ABF92] A. Amir, G. Benson, M. Farach. Alphabet independent two dimensional matching. *Proc. of the 24th ACM Symposium on Theory of Computing*, 1992, pp. 59–68.

[ABF93] A. Amir, G. Benson, M. Farach. Optimal parallel two dimensional pattern matching. *Proc. of the 5th ACM Symposium on Parallel Algorithms and Architectures*, 1993, pp. 79-85.

[ABG92] A. Apostolico, D. Breslauer and Z. Galil. Optimal parallel algorithms for periods, palindromes and squares. *Proc. of the 19th International Colloquium on Automata, Languages and Programming*, Lecture Notes in Computer Science, Vol. 623, 1992, pp. 296–307.

[Ba78] T. J. Baker. A technique for extending rapid exact-match string matching to arrays of more than one dimension. *SIAM Journal on Computing*, Vol. 7, 1978, pp. 533–541.

[Bi77] R. S. Bird. Two dimensional pattern matching. *Information Processing Letters*, Vol. 6, No. 5, 1977, pp. 168–170.

[BG90] D. Breslauer and Z. Galil. An optimal $O(\log\log m)$ time parallel string matching algorithm. *SIAM Journal on Computing*, Vol. 19, 1990, pp. 1051–1058.

[BG91] D. Breslauer and Z. Galil. A lower bound for parallel string matching. *SIAM Journal on Computing*, Vol. 21, 1992, pp. 856–862.

[CGGPR93] M. Crochemore, Z. Galil, L. Gasieniec, K. Park, W. Rytter. Constant-Time Randomized Parallel String Matching. *SIAM Journal on Computing*, 26, 4, 1997, pp. 950-960.

[CGHMR92] M. Crochemore, L. Gasieniec, R. Hariharan, S. Muthukrishnan, W. Rytter. An optimal constant time parallel algorithm for 2D pattern matching. *SIAM Journal on Computing*, 27, 3, , 1998, pp. 668–681.

[CCG+93] R. Cole, M. Crochemore, Z. Galil, L. Gasieniec, R. Hariharan, S. Muthukrishnan, K. Park, W. Rytter. Optimal parallel algorithms for preprocessing and pattern matching in one and two dimensions. *Proc. of the 34rd IEEE Symposium on the Foundations of Computer Science*, 1993, pp. 248–257.

[FRW88] F. Fich, R. Ragde, A. Widgerson. Relations between concurrent-write models of parallel computation, *SIAM Journal on Computing*, Vol. 17, 1988, pp. 606–627.

[GaP] L. Gasieniec, K. Park. Work-Time optimal parallel prefix matching. *Proc. of the European Symposium on Algorithms*, 1994, pp. 471–482.

[GP92] Z. Galil and K. Park. Truly alphabet-independent two dimensional matching. *Proc. of the 33rd IEEE Symposium on the Foundations of Computer Science*, 1992, pp. 247–256.

[GP93] Z. Galil and K. Park. Alphabet-independent two dimensional witness computation. *SIAM Journal on Computing*, 25, 5, 1996, pp. 907-935.

[Ja91] J. JaJa. *Introduction to Parallel Algorithms.* Addison-Wesley, 1991.

[KLP89] Z. Kedem, G. Landau, K. Palem. Optimal parallel suffix-prefix matching algorithm and application. *Proc. of the 1st ACM Symposium on Parallel Algorithms and Architectures*, 1989, pp. 388–398.

[ML84] M.G. Main and R.J. Lorentz. An $O(n\log n)$ algorithm for finding all repetitions in a string. *Journal of Algorithms*, Vol. 5, 1984, pp. 422–432.

[Ra90] P. Ragde. The parallel simplicity of compaction and chaining. *Proc. of the 17th International Colloquium on Automata, Languages and Programming*, Lecture Notes in Computer Science, Vol. 443, 1990, pp. 744–751.

[Vi85] U. Vishkin. Optimal pattern matching in strings. *Information and Control*, Vol. 67, 1985, pp. 91–113.

[Vi90] U. Vishkin. Deterministic sampling – A new technique for fast pattern matching. *Proc. of the 22nd ACM Symposium on Theory of Computing*, 1990, pp. 170–180.

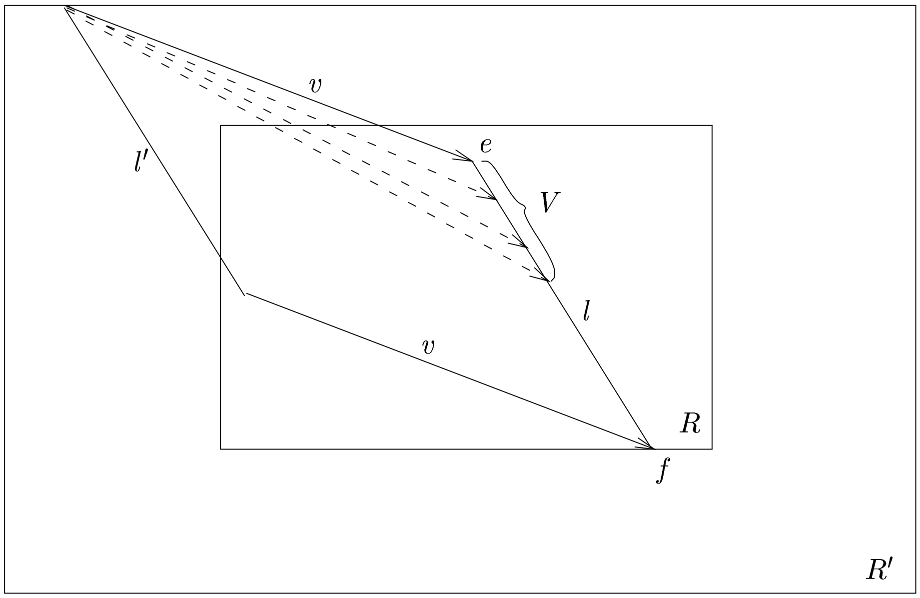

Figure 15: The lines l and l'.

Appendix I

The Procedure Line(V,R,R'). Recall that R, R' are a sub-blocks of p, R is a sub-block of R', and V is some subset of the set of Quad I vectors (Quad II vectors, respectively) with the following property: all the difference vectors of vectors in V are parallel to some Quad I vector (Quad II vector, respectively). V satisfies the additional property that all vectors in V have heads at locations in R when their tails are at the top left corner of R'. $Line(V, R, R')$ finds witnesses with heads in R and tails in R', if any, for vectors in V in $O(\log\log m)$ time and $O(|R|)$ work. We describe only the Quad I case, i.e., the case when vectors in V as well as their difference vectors are Quad I vectors. The Quad II case is similar.

Let v be the vector in V with the smallest row length (with ties broken in favour of the vector with the smallest column length). Recall that all difference vectors of vectors in V are parallel. R is partitioned into lines parallel to these difference vectors. All these lines are processed in parallel.

Consider one such line g and let l be the portion of this line with the property that for all points $x \in l$, $x - v \in R'$ (see Fig.15). Let e be the leftmost topmost point in l. Let f be the rightmost bottommost point in l. Since all vectors in V and their difference vectors are Quad I vectors and since vectors in V have heads in R when their tails are at the topmost leftmost corner of R', any witness for a vector in V with head at a character in g and tail at a character in R' has its head between e and f (both inclusive) in l (see Fig.15). Let l' denote the line in R' such that $x \in l'$ if and only if $x + v \in l$. Clearly, $|l| = |l'|$. Witnesses for vectors in V with heads at locations in l (i.e., tails at locations in l') are found in $O(\log\log m)$ time and $O(|l| + |l'|)$ work. The total work done over all lines g in R is clearly $O(|R|)$.

Let s_1 denote the string comprising, from left to right, the characters in l' between (and including) $e-v$ and $f-v$. Let s_2 denote the string comprising, from left to right, the characters in l between (and including) e and f. Note that both s_1 and s_2 have the same length. Let s be the string obtained by concatenating s_1 and s_2. Let the characters in s be indexed from 0 to $|s|-1$. Clearly, each character x between e and f in l and between $e-v$ and $f-v$ in l' is associated with an unique character in s; let i_x denote the index of this character in s.

Using the string witness computation algorithm [BG90, ABG92], for each i, $|s_1| \leq i \leq |s|-1$, an index j, if any, such that $s[j] \not\equiv s[j-i]$ is found in $O(\log\log m)$ time and $O(|s|)$ work. We call j a *witness-index* of i. It can be easily seen that (c,d), $c \in l', d \in l$, is a witness for w if and only if i_d is a witness-index of $|s_1| + rl(w) - rl(v)$.

Substep B.1. Consider the set S of valid Quad I vectors which fall in the same row r. We show how Substep B.1 is performed for this row in $O(\log\log m)$ time and $O(m)$ work.

We associate each vector in S with the character upon which its head falls when its tail is at the top left corner of p'; such a character is called a *source*. A source *survives* if a witness has not been found for the vector to which it corresponds; otherwise, it is said to have been *eliminated*.

We describe the algorithm for Step B next. In this algorithm, row r will be partitioned into subrows; a subrow is said to have been processed if all surviving sources in that subrow are consistent, i.e., if sources a, b in that subrow survive then the vector $b-a$ (the vector joining a to b) does not have a witness in p'.

First, r is partitioned into disjoint subrows of size $e = \lceil \log\log m \rceil$. Each subrow is processed sequentially by a single processor in $O(\log\log m)$ time. Consider a particular subrow sr. The sources $s_1, s_2, \ldots, s_k$ (in order from left to right) in sr are considered sequentially in order, where $k \leq e$. When source s_i is reached, all surviving sources to the left of s_i are consistent. The vector associated with s_i is then duelled with each of the vectors associated with the sources to the left of s_i which still survive in right to left order; this is done until either s_i is eliminated or all sources to the left of s_i are eliminated or s_i is found to be consistent with some surviving source to its left. At this point, either s_i has been eliminated or all surviving sources to the left of s_i are consistent with s_i. s_{i+1} is considered next. This process continues until s_k has been considered.

Next, a recursive divide-and-conquer algorithm is used. There are $O(\log\log m)$ levels of recursion; each level of recursion takes $O(1)$ time and does $O(\frac{m}{e}) = O(\frac{m}{\log\log m})$ work. At the ith level of recursion, $1 \leq i \leq \log\log m$, r is partitioned into disjoint subrows of size at most $\lceil (\frac{m}{e})^{\frac{1}{2^{i-1}}} \rceil \times e$; these subrows are called i-subrows. The size of the subrows at the deepest level of recursion is at most e.

We maintain the invariant that after the ith level of recursion, all sources in each i-subrow are consistent. Consider the ith level of recursion. If this is the deepest level of recursion, i.e., i-subrows have size at most e, then all sources in each i-subrow are consistent. Assume that this is not the deepest level of recursion. Let sr be an i-subrow. All pairs of distinct $(i+1)$-subrows which overlap sr are processed in parallel in $O(1)$ time and $O(1)$ work per pair. Since there are $O(\frac{m}{(\frac{m}{e})^{\frac{1}{2^{i-1}}} \times e})$ i-subrows and each i-subrow has $O((\frac{m}{e})^{\frac{1}{2^i}})$ $(i+1)$-subrows, the total work done in the ith level of recursion is $O(\frac{m}{e})$ as claimed.

It remains to describe how sources in a pair of distinct $(i+1)$-subrows in sr are made consistent in $O(1)$ time and work. If both subrows have at most one surviving source, then the vectors associated with the two sources are duelled. Otherwise, suppose that at least one of the subrows has two surviving sources. Then p' must have a valid horizontal period vector,

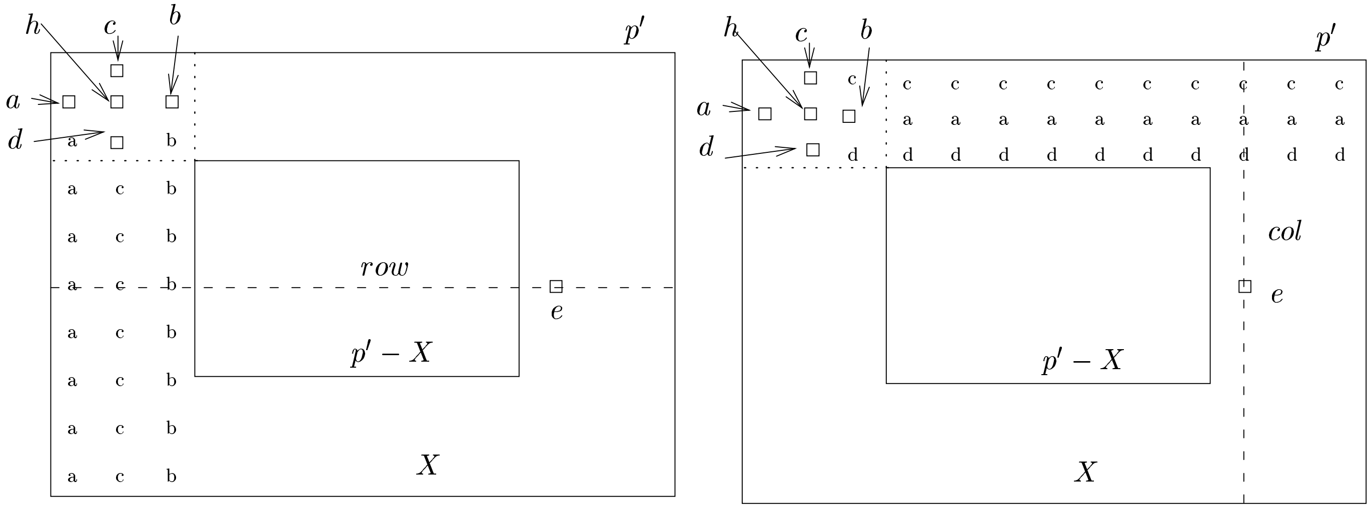

Figure 16: Col-marks and row-marks in Step 1a when X is a fringe; a, b, c, d are sources.

v say; therefore p' is r-oriented. The vector v_a associated with an arbitrary source a in the first subrow is duelled with the vector v_b associated with an arbitrary source b in the second subrow as follows.

If $v_b - v_a$ has no witness in p' then all sources in the first subrow are consistent will all sources in the second subrow. Otherwise, suppose $v_b - v_a$ has a witness in p'. By Lemma 3.1, there is a witness (e, f) for $v_b - v_a$ in the middle $\frac{1}{4}$th fraction of the columns of p'. Recall that the vectors in consideration have row length less than $\frac{m'_1}{8}$ and column length less than $\frac{m'_2}{8}$. This implies that if $e + v_b \in p'$ then $e + v \in p'$ for all vectors v associated with surviving sources in the two subrows, and likewise for $e - v_a$. In addition, this implies that either $e - v_a$ or $e + v_b$ is in p'. Witness (e, f) is used to duel v_a and v_b by comparing e with either $e - v_a$ or $e + v_b$, whichever is in p'. Suppose that $e - v_a \in p'$. If $e - v_a \not\equiv e$ then $(e - v, e)$ is a witness in p' for all vectors v associated with surviving sources in the first subrow. Otherwise, if $e - v_a \equiv e$ then $e \equiv e - v_a \equiv f - v_b \not\equiv f$ and therefore $(f - v, f)$ is a witness in p' for all vectors v associated with surviving sources in the second subrow. The case when $e + v_b \in p'$ is handled similarly.

The Procedure Verify(X,Y,V). Recall that X and Y are both identically sized rectangular sub-blocks of p or fringes of p' with $X = Y$, and V is a subset of the locations in X with the following property: for all $a, b \in V$, copies of Y placed with the top left corner at a and b match each other wherever they both overlap X. In addition, when X is a fringe of p', all locations in V are strictly above and to the left of the region $p' - X$.

Let Y_a denote the copy of Y placed with top left corner at a. For each $a \in V$, this procedure determines a location $bad_a \in X$, if any, where Y_a mismatches X. We show how this procedure is performed in $O(1)$ time and $O(|X|)$ work.

A *source* is defined to be a location in V. Imagine copies of the Y placed with top left corners at the various sources.

There are three steps, Steps 1–3.

Step 1. In this step, each location e in X is marked with a source a, if any, with the property that Y_a overlaps e. There are two stages in this step. Lemma 9.15 shows the correctness of this step.

Step 1.a. Two kinds of marking, called *row-marking* and *col-marking*, are performed in this step (see Fig.16). Each location in X which has a source on or above it in the same column is col-marked with the uppermost source in its column. Similarly, each location in X which has a source on or to its left in the same row is row-marked with the leftmost source in its row. This is accomplished easily in $O(1)$ time and $O(|X|)$ work using the algorithm of [FRW88].

Step 1.b. In each row of X, the leftmost and rightmost col-marked locations are determined in $O(1)$ time and $O(|X|)$ work using the algorithm of [FRW88] for each row. In each column of X, the topmost and bottommost row-marked locations are determined in $O(1)$ time and $O(|X|)$ work.

Step 1.c. Each location $e \in X$ is processed in parallel in $O(1)$ time and work as follows. Let e be in row row and column col (see Fig.16). Let a and b be the sources with which the leftmost and rightmost col-marked locations in row, if any, are col-marked. Let c and d be the sources with which the topmost and bottommost row-marked locations in col, if any, are row-marked. If e is overlapped by one Y_a, Y_b, Y_c, Y_d, then it is marked with the corresponding source; otherwise, it is left unmarked.

Lemma 9.15 *If location $e \in X$ is marked with source g in Step 1.c then Y_g overlaps e. If e is left unmarked in Step 1.c then for no source $g \in V$ does Y_g overlap e.*

Proof. The first part of the lemma is obvious. Consider the second part. Suppose there is a source $g \in V$ such that Y_g overlaps e. We show that e will be marked in this case. Let e be in row row and column col of X.

Let a and b be the sources with which the leftmost and rightmost col-marked locations in row, if any, are col-marked (see Fig.16). To see that a and b are defined though not necessarily distinct, consider source g. As all sources are above and to the left of $p' - X$, all locations below g in p' which are in the same column as g are col-marked, though not necessarily with g. In particular, the element in row row in the same column as g is col-marked and therefore a and b are defined, though not necessarily distinct.

Clearly, g is in or between the columns containing a and b. Further, e is in the same column as g or to its right. Since some location in row row is col-marked with a, a is in or above row row. Therefore, e is in the same row as a or below it and in the same column as a or to its right.

It follows that if X is a sub-block of p then Y_a overlaps e and therefore e is marked in Step 1.c.

Next, suppose X is a fringe of p'. Let p'_f denote the copy of p' with top left corner at location $f \in X$. Let e_f denote the location in p'_f, if any, which overlaps e. From the above paragraph, it follows that e_a is defined. Since Y_g overlaps e, e_g is defined and $e_g \notin p'_g - Y_g$. There are two cases to be considered next, depending upon whether e_g is to the right or left of $p'_g - Y_g$. The cases when e_g is above or below $p'_g - Y_g$ are shown similarly. We show that either Y_a or Y_b overlaps e in each case and therefore, e is marked in Step 1.c.

First, suppose e_g is to the right of $p'_g - Y_g$. Since a is to the left of e, e_a is to the right of $p'_a - Y_a$. Therefore, $e_a \notin p'_a - Y_a$. Since e_a is defined, Y_a overlaps e in this case.

Second, suppose e_g is to the left of $p'_g - Y_g$. There are two subcases. First, suppose e is either to the left of or above the top left corner h of the region $p' - X$. Clearly, $p'_a - Y_a$ is entirely below and to the right of h and therefore $e_a \notin p'_a - Y_a$. Since e_a is defined, Y_a overlaps e in this case. Second, suppose e is below and to the right of h. Then since b is above and to the left of e, e_b is defined. Since b is aligned with or to the right of g and since e_g is to the left of $p'_g - Y_g$, e_b is to the left of $p'_b - Y_b$. Therefore, Y_b overlaps e in this case. □

Step 2. Each character e in X marked with some source a is compared with the character in Y_a which overlaps e.

A location in X which mismatches in Step 2 is called a *bad* location.

Step 3. In this step, each source e is marked with a bad location, if any, overlapped by Y_e. Thus, after Step 3, for every source e, a location $bad_e \in X$, if any, such that the character overlapping bad_e in Y_e differs from the character at bad_e is determined. This step is similar to Step 1. All marks made in Step 1 are erased before Step 3 begins.

Step 3.a. Two kinds of marking, called *row-marking* and *col-marking*, are performed in this step. Each location in X with a bad location on or below it in the same column is col-marked with the lowermost bad location in its column. Similarly, each location in X with a bad location on or to its right in the same row is row-marked with the rightmost bad location in its row. This is accomplished easily in $O(1)$ time and $O(|X|)$ work.

Step 3.b. In each row of X, the leftmost and rightmost col-marked locations are determined in $O(1)$ time and $O(|X|)$ work. In each column of X, the topmost and bottommost row-marked locations are determined in $O(1)$ time and $O(|X|)$ work.

Step 3.c. Each source $e \in V$ is processed in parallel in $O(1)$ time and work as follows. Let e be in row row and column col. Let a and b be the bad locations with which the leftmost and rightmost col-marked locations in row, if any, are col-marked. Let c and d be the bad locations with which the topmost and bottommost row-marked locations in col, if any, are row-marked. If Y_e overlaps one of a, b, c, d then it is marked with the corresponding bad character, and otherwise, it is left unmarked.

Lemma 9.16 states the correctness of Step 3. Its proof is similar to the proof of Lemma 9.15.

Lemma 9.16 *If source $e \in V$ is marked with bad location $g \in X$ in Step 3.c then Y_e overlaps g. If e is left unmarked in Step 3.c then for no bad location $g \in X$ does Y_e overlap g.*

Appendix II: A New Periodicity Property

In this section, we describe a new periodicity property of two dimensional patterns. This property appeared implicitly in Lemmas 3.4 and Lemma 9.12 and played a key role in the algorithm described in this paper. Since this property is of independent interest and may well be useful in other algorithms, we describe it explicitly here.

Definitions. In this section we will consider only vectors which have row length less than $\frac{m_1}{8}$ and column length less than $\frac{m_2}{8}$. The term *valid vector* is redefined to denote vectors whose lengths are constrained as above. For the purpose of this section, a *source* is defined to be a point in p on which the head of some valid period vector of p lies when its tail is at the top left corner of p.

Amir and Benson [AB92] introduced the following classification which was subsequently refined by Galil and Park [GP92, GP93]. p is classified into one of the following four categories, depending upon the nature of its period vectors.

Non-Periodic: p has no valid period vectors.
Lattice-Periodic: p has a valid Quad I and a valid Quad II period vector.
Radiant-Periodic: All valid period vectors of p are either Quad I vectors or Quad II vectors; further, some two valid period vectors are non-parallel.
Line-Periodic: All valid period vectors of p are either Quad I vectors or Quad II vectors; further they are all parallel.

The new property we obtain concerns the distribution of sources in the Radiant-Periodic case. Suppose p is Radiant-Periodic. Without loss of generality, assume that all valid period vectors are Quad I vectors.

Definitions. Let w_1, w_2 denote any two non-parallel valid period vectors of p. Let U denote the top left corner sub-block of p of size $\lceil \frac{m_1}{2} \rceil \times \lceil \frac{m_2}{2} \rceil$. Let A be the top left corner sub-block of p of size $\lceil \frac{m_1}{8} \rceil \times \lceil \frac{m_2}{8} \rceil$. Clearly, all sources are in A. For any point x in A, let $TL(x)$ and $BR(x)$ be maximal sub-blocks of A defined as follows. Points in $TL(x)$ are vertically aligned with or above x and horizontally aligned with or to the left of x. Points in $BR(x)$ are strictly below and to the right of x. Note that if x is the bottom left corner of A then $BR(x)$ is empty.

The following property follows from Lemma 8 in [GP93].

Lemma 9.17 [Galil,Park [GP93]] *There exists a valid Quad I vector v_1 and a valid Quad II vector v_2 satisfying the following properties. v_1 and v_2 are period vectors of U. In addition, a valid vector is a period vector of U if and only if it is a linear combination of v_1, v_2. It follows that point $x \in A$ is a source only if x is a (v_1, v_2)-lattice point with respect to the top left corner of p.*

We show the following lemma.

Lemma 9.18 ["Cat with a Broken Tail" Lemma] *There exist points x, y (not necessarily distinct) in A with the following properties (see Fig.15[3]).*
(1) $y - x$ is a Quad I vector.
(2) All sources are in $TL(x) \cup BR(y)$.

[3]Dotted lines in the boundary of p signify the fact that the figure is not to scale.

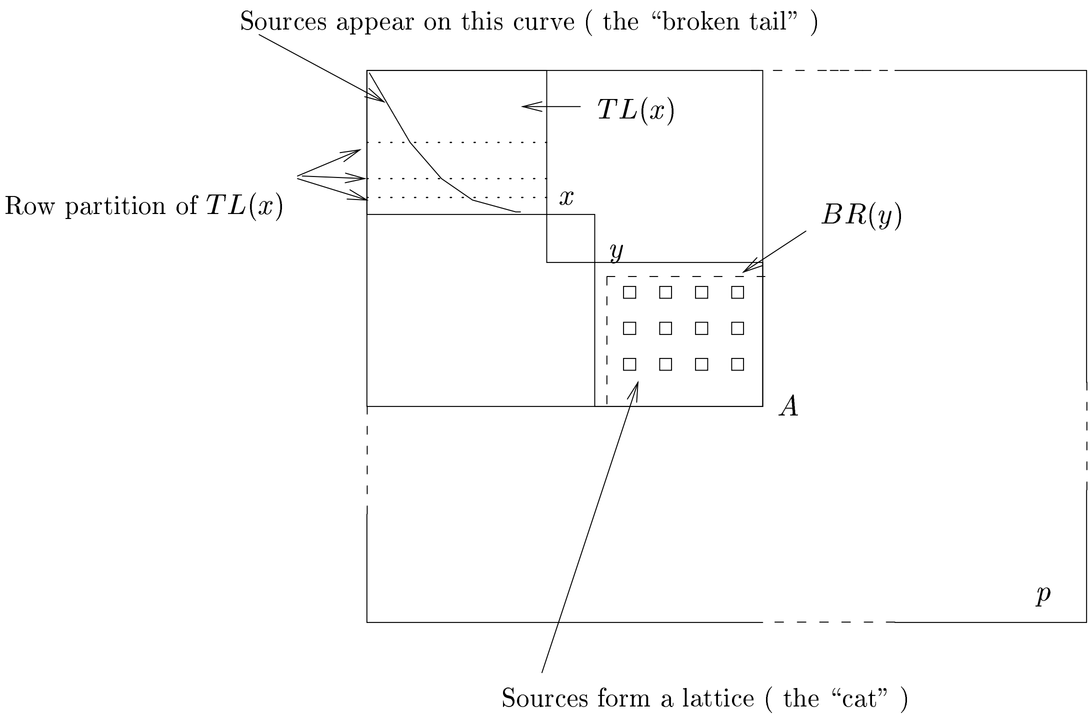

Figure 17: The "cat" with the "broken tail".

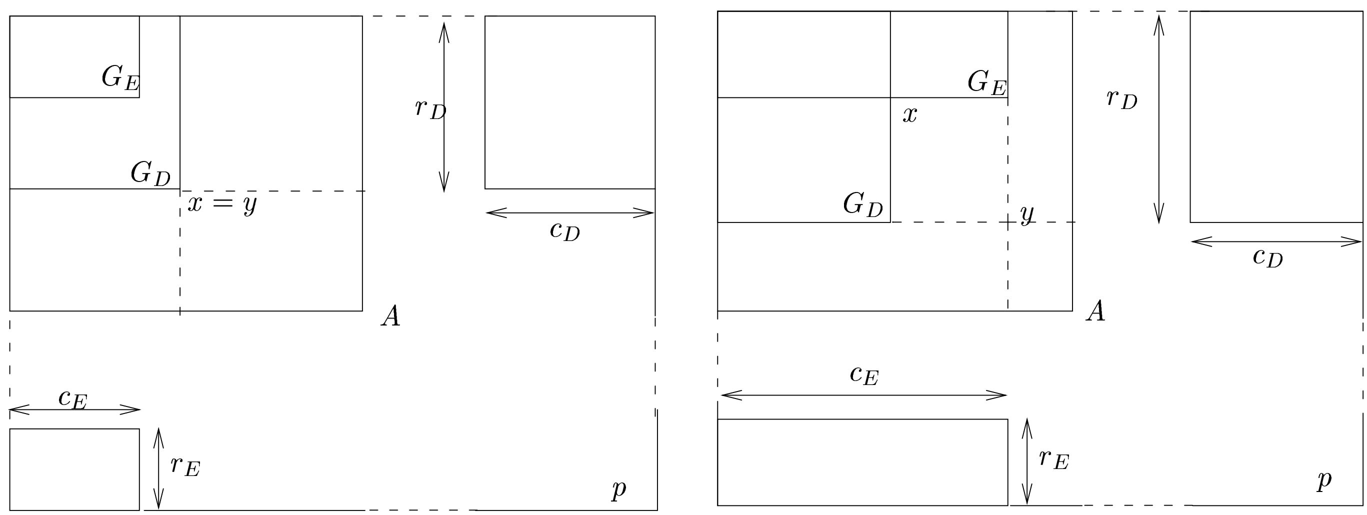

a) G_E contained in G_D

b) Neither of G_D, G_E contained in the other

Figure 18: G_D, G_E, x and y.

(3) Point z in $BR(y)$ is a source if and only if z is a (v_1, v_2)-lattice point with respect to the top left corner of p.
(4) Sources in $TL(x)$, if any, appear on a piecewise linear curve consisting of $O(\log m)$ straight line segments. More precisely, one of the following is true:
(a) all sources which lie in the same segment of the row partition of the rows of p intersecting $TL(x)$ are collinear (recall the definition of row partition from Section 5).
(b) all sources which lie in the same segment of the column partition of the columns of p intersecting $TL(x)$ are collinear (with column partition defined analogous to row partition).

The rest of this section is devoted to the proof of Lemma 9.18.

Definitions. Let D be the $\lceil \frac{m_1}{4} \rceil \times \lceil \frac{m_2}{4} \rceil$ sized sub-block of p with the same top right corner as p. Let E be the $\lceil \frac{m_1}{4} \rceil \times \lceil \frac{m_2}{4} \rceil$ sized sub-block of p with the same bottom left corner as p. Consider the (v_1, v_2)-lattice with respect to the center e of the leftmost column in U (with ties for the center broken in favour of the lower candidate) and let C denote the cell bounded by $e, e+v_1, e+v_2, e+v_1+v_2$.

Lemma 9.19 follows from Lemma 12 in [GP93].

Lemma 9.19 [Galil,Park [GP93]] *All defects in p with respect to C are in D and E.*

Lemma 9.20 *Let v be a valid vector which is a linear combination of v_1, v_2. Let z, z' be points in p such that $z' = z + v$. If exactly one of z, z' is a defect with respect to C then $z \not\equiv z'$, i.e., (z, z') is a witness for v. If neither z nor z' is a defect with respect to C then $z \equiv z$, i.e., (z, z') is not a witness for v.*

Proof. Consider U', the maximal sub-block of $p - D - E$ with the same top left corner as p. v must be a period vector of U', because if v has a witness (c, d) in U' then either c or d must be a defect with respect to C, a contradiction to Lemma 9.19. Clearly, for all points $x \in C$, $x, x + v \in U'$. By Lemma 9.19, it follows that C is safe for v in U'. The lemma follows from Corollary 6.6 with U', p, z, z' replacing A, A', x, y, respectively. □

Corollary 9.21 *Either D or E has a defect with respect to C.*

Proof. By Lemma 9.20 and the fact that v_2 is not a period vector of p (recall that all valid period vectors of p are Quad I vectors), p has a defect with respect to C. The corollary follows from Lemma 9.19. □

Definitions. Let r_D, c_D be the smallest numbers such that all defects in D are contained in the top r_D rows and rightmost c_D columns of D. Note that $r_D = c_D = 0$ if there are no defects in D. Similarly, let r_E, c_E be the smallest numbers such that all defects in E are contained in the bottommost r_E rows and leftmost c_E columns of D. Note that by Corollary 9.21, if $r_D = c_D = 0$ then $r_E \neq 0, c_E \neq 0$ and vice versa.

Let G_D be the top left corner sub-block of p of size $r_D \times c_D$. Let G_E be the top left corner sub-block of p of size $r_E \times c_E$ (see Fig.16[4]). Let H be the smallest top left corner sub-block of p which encloses both G_D, G_E. Let H' be the largest top left corner sub-block of p containing only points in $G_D \cup G_E$. Recall Lemma 9.18. y is defined to be the bottom right corner of $H \cap A$. x is defined to be the bottom right corner of $H' \cap A$. Clearly, $y - x$ is a Quad I vector and the first part of Lemma 9.18 follows.

Lemma 9.22 along with Lemma 9.17 shows the third part of Lemma 9.18. Lemma 9.23 shows the second part of Lemma 9.18. To show the fourth part of Lemma 9.18, assume without loss of generality that $r_D \geq r_E$ (cases (a) and (b) in Fig.16). Then Lemma 9.24 shows part 4a of Lemma 9.18. If $r_D < r_E$, a proof analogous to the one in Lemma 9.24 shows part 4b of Lemma 9.18.

Definitions. We redefine the term *fall* as follows in this section. We say that a vector v *falls* in some sub-block of p if its head is in that sub-block when its tail is at the top left corner of p.

Lemma 9.22 *Suppose valid vector v is a period vector of U and falls in $BR(y)$. Then v is a period vector of p.*

Proof. Since v falls in $BR(y)$, v is a Quad I vector. Further, $rl(v) \geq \max\{r_D, r_E\}$ and $cl(v) \geq \max\{c_D, c_E\}$. By Lemma 9.17, v is a linear combination of v_1, v_2. When one endpoint of v is at a defect in either D or E, the other endpoint is outside p. From Lemma 9.20, it follows that v has no witnesses in p. □

Lemma 9.23 *No valid period vector of p falls in $A - TL(x) - BR(y)$.*

Proof. Consider a valid period vector v of p. Suppose for a contradiction that v falls in $A - TL(x) - BR(y)$. Then one of the following is true.

1. $rl(v) < r_D$ and $cl(v) \geq c_D$.

[4]Dotted lines in the boundary of p signify the fact that the figure is not to scale.

2. $rl(v) \geq r_D$ and $cl(v) < c_D$.

3. $rl(v) < r_E$ and $cl(v) \geq c_E$.

4. $rl(v) \geq r_E$ and $cl(v) < c_E$.

We consider the first case. The other cases are handled similarly.

Recall that v is a Quad I vector and that, by Lemma 9.17, it is a linear combination of v_1, v_2. Since $rl(v) \geq 0$, $r_D > 0$ and therefore, D has a defect. Let z be a b-defect in D. Since $rl(v) < r_D$ and $cl(v) \geq c_D$, $z - v$ is in p and a non-defect. By Lemma 9.20, $(z - v, z)$ is a witness for v, a contradiction. □

Lemma 9.24 *Suppose that $r_D \geq r_E$. Consider the row partition of those rows of p which intersect with $TL(x)$. Let u, v, w be valid period vectors of p. Suppose, without loss of generality that*
(a) $rl(w) \geq rl(v) \geq rl(u)$,
(b) if $rl(w) = rl(v)$ then $cl(w) > cl(v)$, and
(c) if $rl(v) = rl(u)$ then $cl(v) > cl(u)$.
If u, v, w fall in the jth segment of the this row partition then $w - v$ and $v - u$ are parallel Quad I vectors.

Proof. Since $r_D \geq r_E$ and since one of D, E contains a defect with respect to C, $r_D > 0$ and D contains a defect with respect to C. It follows from the definition of r_D that the r_Dth row of p contains a defect with respect to C in D; let z be one such defect.

Let k be the number of rows in the jth segment. Let sub-block B denote the intersection of the bottommost k of the top r_D rows of p and the rightmost $\lceil \frac{7m_2}{8} \rceil$ columns of p. Clearly, B is completely below the jth segment and z lies in B. Further, $z - w, z - v, z - u$ are all in p.

First, we claim that $w - v$ and $v - u$ are Quad I vectors. We show this for $w - v$; the proof for $v - u$ is similar. By Lemma 9.17, w, v and therefore $w - v$ are linear combinations of v_1, v_2. If $w - v$ is not a Quad I vector (i.e., $v - w$ is a non-horizontal non-vertical Quad II vector) then $z + (w - v)$ is a non-defect with respect to C. By Lemma 9.20, $(z, z + w - v)$ is a witness for $w - v$. Then either $(z, z - v)$ is a witness for v or $(z - v, z + w - v)$ is a witness for w, a contradiction.

Second, we claim that $w - v$ and $v - u$ are period vectors of B. Suppose for a contradiction that $w - v$ is not a period vector of B. Let (c, d) be a witness for $w - v$ in B. Since w, v are Quad I vectors which fall in the jth segment and B does not intersect with the leftmost $\lfloor \frac{m_2}{8} \rfloor$ columns of p, $c - v$ is in p and therefore, either $(c - v, c)$ is a witness for v or $(d - w, d)$ is a witness for w, a contradiction. A similar proof shows that $v - u$ is a period vector of B.

Finally, suppose for a contradiction that $w - v$ and $v - u$ are not parallel. Note that B has at least $rl(w - v) + rl(v - u) + 1$ rows and $\lceil \frac{7m_2}{8} \rceil$ columns. Further, z is in the bottommost row of B but not in the leftmost $cl(w - v) + cl(v - u) - 1$ columns of B; it follows that z is in $good_B(w - v, v - u)$. Consider the $(w - v, v - u)$-lattice with respect to z. The sizes of $B, D, w - v, v - u$ imply that there is a lattice point b in $B - D$ which is in $good_B(w - v, v - u)$. By Lemma 6.1, there is a $(w - v, v - u)$-lattice path completely contained in B between z and b. Since $w - v, v - u$ are linear combinations of v_1, v_2 and period vectors of B, it follows that b is a defect with respect to C in $p - D - E$, a contradiction to Lemma 9.19. □

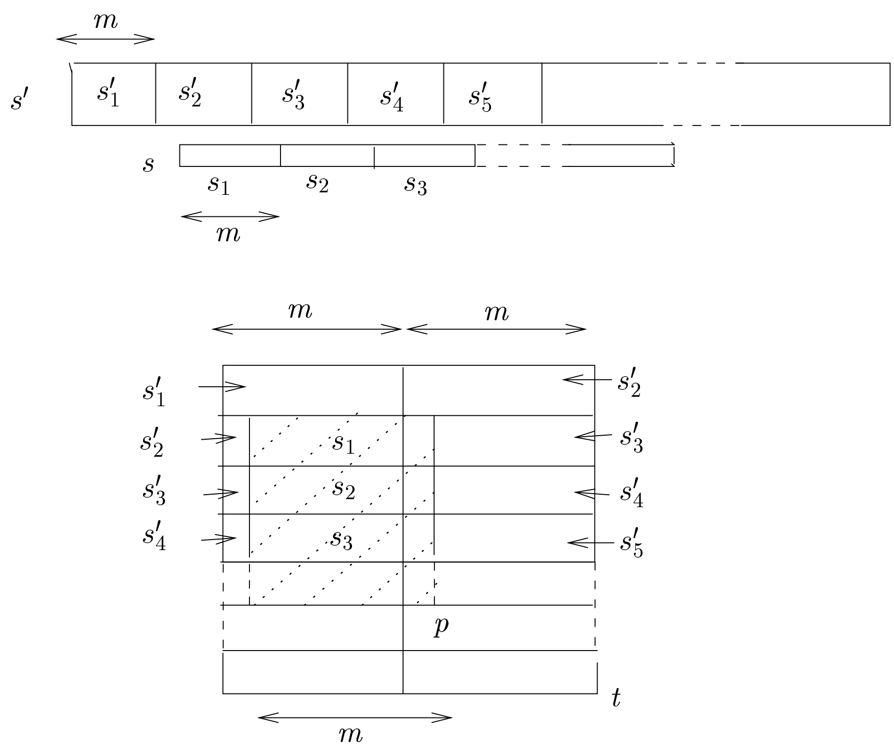

Figure 19: s, s' and p, t.

Appendix III: A lower bound for matching squares

We give a simple many-one reduction from string matching to the problem of matching a square pattern against a square text. More precisely, we reduce the problem of matching a length m^2 pattern string against a length $2m^2-1$ text string to the problem of matching a size $m \times m$ 2D pattern against a size $2m \times 2m$ 2D text. It follows from [BG91] that the latter problem requires $\Omega(\log\log m)$ rounds of character comparisons with $O(m^2)$ processors.

Let s be the pattern string of length m^2 and let s' be the text string of length $2m^2-1$. Let the characters in each string be indexed from 0 onwards. A 2D pattern p of size $m \times m$ is defined as follows: the ith row of p comprises the characters $im \dots im+m-1$ of s, for $i = 0 \dots m-1$. A 2D text t of size $2m \times 2m$ is defined as follows: the ith row of t comprises the characters $im \dots im+2m-1$ of s' for $i = 0 \dots 2m-2$, and the bottommost row of t comprises the characters $(2m-1)m \dots (2m-1)m+m-2$ of s' followed by some $m+1$ arbitrary characters.

Clearly, there is an occurrence of pattern s beginning at character i in s' if and only if there is an occurrence of pattern p with top left corner at character j in row j' of t, where $0 \le i \le m^2-1$, $j' = \lfloor \frac{i}{m} \rfloor$, and $j = i \bmod m$ (see Fig.17). This completes the reduction.