

A Linear Time and Space Algorithm for Detecting Path Intersection*

Srećko Brlek¹, Michel Koskas², and Xavier Provençal³

¹ Laboratoire de Combinatoire et d'Informatique Mathématique,
Université du Québec à Montréal,
CP 8888 Succ. Centre-ville, Montréal (QC) Canada H3C 3P8
brlek.srecko@uqam.ca

² UMR AgroParisTech/INRA 518, 16 rue Claude Bernard, 75231 Paris Cedex 05
michel.koskas@agroparistech.fr

³ LAMA, Université de Savoie, 73376 Le Bourget du Lac, France
LIRMM, Université Montpellier II, 161 rue Ada, 34392 Montpellier, France
xavier.provençal@lirmm.fr

Abstract. For discrete sets coded by the Freeman chain describing their contour, several linear algorithms have been designed for determining their shape properties. Most of them are based on the assumption that the boundary word forms a closed and non-intersecting discrete curve. In this article, we provide a linear time and space algorithm for deciding whether a path on a square lattice intersects itself. This work removes a drawback by determining efficiently whether a given path forms the contour of a discrete figure. This is achieved by using a radix tree structure over a quadtree, where nodes are the visited grid points, enriched with neighborhood links that are essential for obtaining linearity.

Keywords: Freeman code, lattice paths, self-intersection, radix tree, discrete figures, data structure.

1 Introduction

Many problems in discrete geometry involve the analysis of the contour of discrete sets. A convenient way to represent them is to use the well-known Freeman chain code [1,2] which encodes the contour by a word w on the four letter alphabet $\Sigma = \{a, b, \bar{a}, \bar{b}\}$, corresponding to the unit displacements in the four directions (right, up, left, down) on a square grid. Among the many problems that have been considered in the literature, we mention : computations of statistics such as area, moment of inertia [3,4], digital convexity [5,6,7], and tiling of the plane by translation [8,9]. All of these problems are solved by using algorithms shown to be linear in the length of the contour word, but often it is assumed that the path encoded by this word does not intersect itself. While it is easy to check that a word encodes a closed path (by checking that the word contains as many a as \bar{a} , and as many b as \bar{b}), checking that it does not intersect

* With the support of NSERC (Canada).

itself requires more work. The problem amounts to check if a grid point is visited twice. Of course, one might easily provide an $\mathcal{O}(n \log n)$ algorithm where sorting is involved, or use hash tables providing a linear time algorithm on average but not in worst case.

The goal of this paper is to remove this major drawback. Indeed, we provide a linear time and space algorithm in the worst case checking if a path encoded by a word visits any of the grid points twice. Section 2 provides the basic definitions and notation used in this paper. It also contains the description of the data structures used in our algorithm: it is based on a quadtree structure [10], used in a novel way for describing points in the plane, combined with a radix tree (see for instance [11]) structure for the labels. In Section 3 the algorithm is described in details. The time and space complexity of the algorithm is carried out in Section 4, followed by a discussion on complexity issues, with respect to the size of numbers and bit operations involved. Finally a list of possible applications is provided showing its usefulness.

2 Preliminaries

A word w is a finite sequence of letters $w_1 w_2 \dots w_n$ on a finite alphabet Σ , that is a function $w : [1..n] \rightarrow \Sigma$, and $|w| = n$ is its *length*. Therefore the i th letter of a word w is denoted w_i , and sometimes $w[i]$ when we emphasize the algorithmic point of view. The empty word is denoted ε . The set of words of length n is denoted Σ^n , that of length at most n is $\Sigma^{\leq n}$, and the set of all finite words is Σ^* , the free monoid on Σ . Similarly, the number of occurrences of the letter $\alpha \in \Sigma$, is denoted $|w|_\alpha$. From now on, the alphabet is fixed to $\Sigma = \{a, b, \bar{a}, \bar{b}\}$. To any word $w \in \Sigma^*$ is associated a vector \vec{w} by the morphism $\vec{\cdot} : \Sigma^* \rightarrow \mathbb{Z} \times \mathbb{Z}$ defined on the elementary translation $\vec{\epsilon} = (\epsilon_1, \epsilon_2)$ corresponding to each letter $\epsilon \in \Sigma$:

$$\vec{a} = (1, 0), \vec{\bar{a}} = (-1, 0), \vec{b} = (0, 1), \vec{\bar{b}} = (0, -1),$$

and such that $\vec{u \cdot v} = \vec{u} + \vec{v}$. For sake of simplicity we often use the notation \mathbf{u} for vectors \vec{u} .

The set of elementary translations allows to draw each word as a 4-connected path in the plane starting from the origin, going *right* for a letter a , *left* for a letter \bar{a} , *up* for a letter b and *down* for a letter \bar{b} . This coding proved to be very convenient for encoding the boundary of discrete sets and is well known in discrete geometry as the *Freeman chain code* [1,2]. It has been extensively used in many applications and allowed the design of elegant and efficient algorithms for describing geometric properties.




Fig. 1. Path encoded by the word $w = aaab\bar{b}baab\bar{a}$

design of elegant and efficient

The underlying principle of our algorithm is to build a graph whose nodes represent points of the plane. For that purpose, the plane is partitioned as in Fig. 2, where the point (2,1) is outlined with its four sons (solid arrows) and its four neighbors (dashed arrows). The sons of a node are grouped in grey zones, while dashed lines separate the levels of the tree. Each node has two possible states : *visited* or *not visited*. New nodes are created while reading the word $w = w_1w_2 \cdots w_n$ from left to right. For each letter w_i , the node corresponding to the point $w[1..i]$ is written as *visited* and of course, if at some point a node is visited twice then the path is not self-avoiding and the algorithm stops. During the process is built a graph $\mathcal{G} = (N, R, T)$ where N is a set of nodes associated to points of the plane, R and T are two distincts sets of oriented edges. The edges in R give a quadtree structure on the nodes while the edges in T are links from each node to its *neighbors*, for which we give a precise definition.

Definition 1. *Given a point $(x, y) \in \mathbb{Z}^2$, we say that (x', y') is a neighbor of (x, y) if there exists $\epsilon \in \Sigma$ such that $(x', y') = (x, y) + \epsilon = (x + \epsilon_1, y + \epsilon_2)$.*

When we want to discriminate the neighbors of a given point (x, y) , for each $\epsilon \in \Sigma$, we say that (x', y') is an ϵ -neighbor of (x, y) if $(x', y') = (x, y) + \epsilon$.

2.1 Data Structure

First, we assume that the path is coded by a word w starting at the origin $(0, 0)$, and stays in the first quadrant $\mathbb{N} \times \mathbb{N}$. This means that the coordinates of all points are nonnegative. Subsequently, this solution is modified in order to remove this assumption. Note that in $\mathbb{N} \times \mathbb{N}$, each point has exactly four neighbors with the exception of the origin $(0, 0)$ which admits only two neighbors, namely $(0, 1)$ and $(1, 0)$, and the points on the half lines $(x, 0)$ and $(0, y)$ with $x, y \geq 1$ which admit only three neighbors (see Fig. 2).

Let $\mathbb{B} = \{0, 1\}$ be the base for writing integers. Words in \mathbb{B}^* are conveniently represented in the *radix order* by a complete binary tree (see for instance [11,12]), where the level k contains all the binary words of length k , and the order is given by the breadth-first traversal of the tree. To distinguish a natural number $x \in \mathbb{N}$ from its representation we write $\mathbf{x} \in \mathbb{B}^*$. The edges are defined inductively by the rewriting rule $\mathbf{x} \longrightarrow \mathbf{x} \cdot 0 + \mathbf{x} \cdot 1$, with the convention that 0 and 1 are the labels of, respectively, the left and right edges of the node having value \mathbf{x} . This representation is extended to $\mathbb{B}^* \times \mathbb{B}^*$ as follows.

A quadtree with a radix tree structure for points in the integer plane. As usual, the concatenation is extended to the cartesian product of words by setting for $(\mathbf{x}, \mathbf{y}) \in \mathbb{B}^* \times \mathbb{B}^*$, and $(\alpha, \beta) \in \mathbb{B} \times \mathbb{B}$

$$(\mathbf{x}, \mathbf{y}) \cdot (\alpha, \beta) = (\mathbf{x} \cdot \alpha, \mathbf{y} \cdot \beta).$$




Fig. 2. Partition of $\mathbb{N} \times \mathbb{N}$

Let \mathbf{x} and \mathbf{y} be two binary words having same length. Then the rule

$$(\mathbf{x}, \mathbf{y}) \longrightarrow (\mathbf{x} \cdot 0, \mathbf{y} \cdot 0) + (\mathbf{x} \cdot 0, \mathbf{y} \cdot 1) + (\mathbf{x} \cdot 1, \mathbf{y} \cdot 0) + (\mathbf{x} \cdot 1, \mathbf{y} \cdot 1) \quad (1)$$

defines a $\mathcal{G}' = (N, R)$, sub-graph of $\mathcal{G} = (N, R, T)$, such that :

- (i) the root is labeled $(0, 0)$;
- (ii) each node (except the root) has four sons;
- (iii) if a node is labeled (\mathbf{x}, \mathbf{y}) then $|\mathbf{x}| = |\mathbf{y}|$;
- (iv) edges are undirected, e.g. may be followed in both directions.

By convention, edges leading to the sons are labeled by pairs from the ordered set $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$. These labels equip the quadtree with a *radix tree* structure for Equation (1) implies that $(\mathbf{x}', \mathbf{y}')$ is a son of (\mathbf{x}, \mathbf{y}) , if and only if

$$(\mathbf{x}', \mathbf{y}') = (2\mathbf{x} + \alpha, 2\mathbf{y} + \beta),$$

for some $(\alpha, \beta) \in \mathbb{B} \times \mathbb{B}$. Observe that any pair (x, y) of nonnegative integers is represented exactly once in this tree. Indeed, if $|\mathbf{x}| = |\mathbf{y}|$ (by filling with zeros at the left of the shortest one), the sequence of pairs of digits (the two digits in first place, the two digits in second place, and so on) gives the unique path in the tree leading to this pair. Of course the root may have up to three sons since no edge labeled $(0, 0)$ starts from the root.

Neighboring links. We superpose on G' the neighboring relation given by the edges of T (dashed lines). More precisely, for each elementary translation $\epsilon \in \Sigma$, each node $\mathbb{Z} = (x, y)$ is linked to its ϵ -neighbor $\mathbb{Z} + \epsilon$, when it exists. If a level k is fixed (see Fig. 2), it is easy to construct the graph

$$\mathcal{G}^{(k)} = (N^{(k)}, R^{(k)}, T^{(k)})$$

such that

- (i) if $(\mathbf{x}, \mathbf{y}) \in N^{(k)}$, then $|\mathbf{x}| = |\mathbf{y}| = k$;
- (ii) the functions $N^{(k)} \hookrightarrow \mathbb{N} \times \mathbb{N} \hookrightarrow \mathbb{B}^* \times \mathbb{B}^*$ are injective;
- (iii) $R^{(k)}$ is the radix-tree representation : $(\mathbb{B}^{<k} \times \mathbb{B}^{<k}) \times (\mathbb{B} \times \mathbb{B}) \xrightarrow{\bullet} \mathbb{B}^{\leq k} \times \mathbb{B}^{\leq k}$;
- (iv) the neighboring relation is $T^{(k)} \subseteq N \times (\mathbb{B} \times \mathbb{B}) \times N$.

Note that the labeling in Fig. 3 is superfluous: each node represents indeed an integer unambiguously determined by the path from the root using edges in R ; similarly for the ordered edges. Moreover, if a given subset $M \subset \mathbb{N} \times \mathbb{N}$ has to be represented, then one may trim the unnecessary nodes so that the corresponding graph \mathcal{G}_M is not necessarily complete.




Fig. 3. The point $(2, 1)$ with its neighbors

3 The Algorithm

Adding 1 to an integer $x \in \mathbb{B}^k$ is easily performed by a sequential function. Indeed, every positive integer can be written $x = u1^i0^j$, where $i \geq 1, j \geq 0$, with $u \in \{\varepsilon\} \cup \{\mathbb{B}^{k-i-j-1} \cdot 0\}$. In other words, 1^j is the last run of 1's. The piece of code for adding 1 to an integer written in base 2 is

```

1 : If  $j \neq 0$  then Return  $u1^i0^{j-1}1$ ;
2 :           else If  $u = \varepsilon$  then Return  $1 \cdot 0^i$ ;
3 :           else Return  $u \cdot 0^{-1} \cdot 1 \cdot 0^i$ ;
4 :           end if
5 : end if
    
```


where 0^{-1} means to erase a 0. Clearly, the computation time of this algorithm is proportional to the length of the last run of 1's. Much better is achieved with the radix tree structure, where, given a node \textcircled{z} , its *father* is denoted $f(\textcircled{z})$, and we write $f(x, y)$ or $f(\mathbf{x}, \mathbf{y})$ if its label is (x, y) . The following technical lemma is a direct adaptation to $\mathbb{B}^* \times \mathbb{B}^*$ of the addition above.

Lemma 1. *Let $G^{(k)}$ be the complete graph representing $\mathbb{B}^{\leq k} \times \mathbb{B}^{\leq k}$ for some $k \geq 1, \epsilon \in \Sigma$, and $\textcircled{z} = (\mathbf{x}, \mathbf{y})$ be a node of N^k . If one of the four conditions holds:*

- (i) $\epsilon = a$ and $\mathbf{x}[k] = 0$, (ii) $\epsilon = \bar{a}$ and $\mathbf{x}[k] = 1$,
- (iii) $\epsilon = b$ and $\mathbf{y}[k] = 0$, (iv) $\epsilon = \bar{b}$ and $\mathbf{y}[k] = 1$,

then $f(\textcircled{z}) = f(\textcircled{z} + \epsilon)$. Otherwise, $f(\textcircled{z}) + \epsilon = f(\textcircled{z} + \epsilon)$.

The process is illustrated for case (i) in the diagram on the right where the nodes $(10110, \bullet)$ and $(10111, \bullet)$ share the same father while fathers of neighboring nodes



$(\bullet, 01011)$ and $(\bullet, 01011)$ are distinct but share the same neighboring relation.

Now, assume that the node (\mathbf{x}, \mathbf{y}) exists and that its neighbor $(x + 1, y + 0)$ does not. If $|\mathbf{x}| = |\mathbf{y}| = k$, then the translation $(x, y) + (1, 0)$ is obtained in three steps by the following rules:

1. take the edge in R to $f(\mathbf{x}, \mathbf{y}) = (\mathbf{x}[1..k - 1], \mathbf{y}[1..k - 1])$;
2. take (or create) the edge in T from $f(\mathbf{x}, \mathbf{y})$ to $\textcircled{z} = f(\mathbf{x}, \mathbf{y}) + (1, 0)$;
3. take (or create) the edge in R from \textcircled{z} to $\textcircled{z} \cdot (0, \mathbf{y}[k])$.

By Lemma 1, we have $\textcircled{z} \cdot (0, \mathbf{y}[k]) = (x + 1, y + 0)$, so that it remains to add the neighboring link $(\mathbf{x}, \mathbf{y}) \xrightarrow{a} (x + 1, y + 0)$. Then, a nonempty word $w \in \Sigma^n$ is sequentially processed to build the graph \mathcal{G}_w , and we illustrate the algorithm on the input word $w = aabb$.

- *Initialization*: the algorithm starts with the graph containing only the node $(0,0)$ marked as visited. For convenience, the *non-visited* nodes $(0,1)$, $(1,0)$, and the links from $(0,0)$ to its neighbors are also added. This is justified by the fact that the algorithm applies to nonempty words.




Fig. 4. Initial graph \mathcal{G}_ϵ

Since $(0,0)$ is an ancestor of all nodes, this ensures that every node has an ancestor linked with its neighbors. The *current node* is set to $(0,0)$ and this graph is called the *initial graph* \mathcal{G}_ϵ .

- *Read $w_1 = a$* : this corresponds to the translation $(0,0) + (1,0)$. A neighboring link labeled a starting from $(0,0)$ and leading to the node $(1,0)$ does exist, so the only thing to do is to follow this link and mark the node $(1,0)$ as visited. The current node is now set to $(1,0)$, and this new graph is called \mathcal{G}_a .




Fig. 5. Graph \mathcal{G}_a

- *Read $w_2 = a$* : this time, there is no edge in \mathcal{G}_a labeled a starting from $(1,0)$. Using the translation rules above, we perform:

- (1) go back to the father $f(1,0) = (0,0)$;
- (2) follow the link a to $(1,0)$;
- (3) add node $(2,0) \sim (1,0) \cdot (0,0) = (10,00)$.

Then an edge from $(1,0)$ to $(2,0)$ with label a is added to T . Finally the node $(2,0)$ is marked as *visited*, and becomes the current node.




Fig. 6. Graph \mathcal{G}_{aa}

- *Read $w_3 = b$* : this amounts to perform the translation $(2,0) + (0,1)$. Since the edge to $f(2,0)$ is labeled by $(0,0)$, we know that the second coordinate of the current node $(2,0)$ is even. Therefore, $(2,1)$ and $(2,0)$ must be siblings, that is $f((2,0) + (0,1)) = f((2,0))$. What we need to do then is :

- (1) go back to the father $f(2,0) = (1,0)$;
- (2) follow the edge b if it exists;

Since it does not exist, it must be created to reach the node $(2,1) \sim (10,01) = (1,0) \cdot (0,1)$. Again an edge from $(2,0)$ to $(2,1)$ with label b is added, $(2,1)$ is marked as *visited* and is now the current node.




Fig. 7. Graph \mathcal{G}_{aab}

• *Read* $w_4 = b$: since $f((2, 1))$ has no neighboring link labeled by b , recursion is used to find (or build if necessary) the node corresponding to its translation by b . This leads to the creation of the node $(1, 1) \sim (0, 0) \cdot (1, 1)$ marked as *non-visited*. Then, the node $(2, 2) \sim (1, 1) \cdot (0, 0)$ is added, marked as *visited*, and becomes the current node. Note that a neighboring links between $(1, 0)$ and $(1, 1)$, $(2, 1)$ and $(2, 2)$ are added to avoid eventual searches.




Fig. 8. Graph \mathcal{G}_{aabb}

The algorithm **readWord** sequentially reads $w \in \Sigma^*$, builds dynamically the graph G_w marking the corresponding node as visited, and determines if the path coded by w is self-intersecting, i.e. if some node is visited at least twice.

Algorithm 1 (readWord)
Input: $w \in \{a, b, \bar{a}, \bar{b}\}^*$
 0: $\mathcal{G} \leftarrow \mathcal{G}_\epsilon$; $\textcircled{c} \leftarrow \text{root of } \mathcal{G}$;
 1: **For** i **from** 1 **to** $|w|$ **do**
 2: $\epsilon \leftarrow w_i$;
 3: $\textcircled{z} \leftarrow \text{findNeighbor}(\mathcal{G}, \textcircled{c}, \epsilon)$;
 4: **If** \textcircled{z} is *visited* **then**
 5: w is self-intersecting.
 6: **end if**
 7: **Mark** \textcircled{z} as *visited*;
 8: $\textcircled{c} \leftarrow \textcircled{z}$;
 9: **end for**
 10: w is not self-intersecting.

Algorithm 2 (findNeighbor)
Input: $\mathcal{G} = (N, R, T)$; $\textcircled{c} \in N$;
 $\epsilon \in \{a, b, \bar{a}, \bar{b}\}$;
 1: **If** the link $\textcircled{c} \xrightarrow{-\epsilon} \textcircled{z}$ does not exist **then**
 2: $\textcircled{p} \leftarrow f(\textcircled{c})$
 3: **If** $f(\textcircled{c} + \epsilon) = f(\textcircled{c})$ **then**
 4: $\textcircled{r} \leftarrow \textcircled{p}$;
 5: **else**
 6: $\textcircled{r} \leftarrow \text{findNeighbor}(\mathcal{G}, \textcircled{p}, \epsilon)$;
 7: **end if**
 8: $\textcircled{z} \leftarrow \text{son of } \textcircled{r} \text{ corresponding to } \textcircled{c} + \epsilon$;
 9: Add the neighboring link $\textcircled{c} \xrightarrow{-\epsilon} \textcircled{z}$.
 10: **end if**
 11: **return** \textcircled{z} ;

The algorithm **findNeighbor** finds, and creates if necessary, the ϵ -neighbor of a given node. Thanks to Lemma 1, testing the condition on line 3 is performed in constant time. At line 8, if the node \textcircled{z} does not exist, it is created. Clearly, the time complexity of this algorithm is entirely determined by the recursive call on line 6 since all other operations are performed in constant time. Finally, note that after each call to **findNeighbor** on line 3 of **readWord**, there always exist a neighboring link $\textcircled{c} \xrightarrow{-\epsilon} \textcircled{z}$.

4 Complexity Analysis

The key for analyzing the complexity of this algorithm rests on the fact that each recursive call of Algorithm 2 requires the addition of a neighboring link. This implies that given a node $\textcircled{z} \in N$, when all the neighboring links have been added, there will never be another recursive call on \textcircled{z} . Since a node has at most 4 sons and each of these sons has at most 2 neighbors not sharing the

same father, the number of recursive calls on a single node is bounded by 8. It remains to show that the number of nodes in the graph is proportional to $|w|$.

First, consider the *visited* nodes. For each letter read, exactly one node is marked as *visited*, so that their number is $|w|$. In order to bound the number of *non-visited* nodes, we need a technical lemma. Recall that the father function $f : N \setminus \{(0, 0)\} \rightarrow N$ extends to subset of nodes in the usual way: for $M \subseteq N$, the fathers of M are $f(M) = \{f(\textcircled{s}) \mid \textcircled{s} \in M\}$. Moreover, f can be iterated to get $f^h(M)$, the *ancestors* of rank h of a subset M . Clearly, f is a contraction since $|f(M)| \leq |M|$, and there is a unique ancestor of all nodes, namely the root.

Lemma 2. *Let $M = \{n_1, n_2, n_3, n_4, n_5\} \subset N$ a set of five nodes such that $(n_i, n_{i+1}) \in T$ for $i = 1, 2, 3, 4$, then, $|f(M)| \leq 4$.*

Proof. As shown in Fig. 2, the nodes sharing the same father split the plane in 2×2 squares. As a consequence, at least two of the nodes n_1, n_2, n_3, n_4, n_5 must share the same father, providing the bound $|f(M)| \leq 4$. ■

This allows to bound the number of nodes using the fact that all non-visited nodes are ancestors of visited ones: the only exception is the initialisation step where the non-visited nodes $(0, 1)$ and $(1, 0)$ are created as leaves.

Lemma 3. *Given a word $w \in \Sigma^n$ and the graph $\mathcal{G}_w = (N, R, T)$, the number of nodes in N is in $\mathcal{O}(n)$.*

Proof. Let $N_v \subseteq N$ be the set of visited nodes, and h be the height of the tree (N, R) . It is clear that $N = \bigcup_{0 \leq i \leq h} f^i(N_v)$, and so

$$|N| \leq \sum_{0 \leq i \leq h} |f^i(N_v)|. \tag{2}$$

By construction, the set N_v forms a sequence of nodes such that two consecutive ones are neighbors since they correspond to the path coded by w . Thus, by splitting this sequence of nodes in blocks of length 5, the previous lemma applies, and we have

$$|f(N_v)| \leq 4 \left\lceil \frac{|N_v|}{5} \right\rceil \leq \frac{4}{5} (|N_v| + 4). \tag{3}$$

By Lemma 1, two neighboring nodes either share the same father or have different fathers that are neighbors, so it is for the sets $f(N_v), f^2(N_v), \dots, f^h(N_v)$. So, by combining inequations (2) and (3), the following bound is obtained


$$\begin{aligned} |N| &\leq \sum_{0 \leq i \leq h} |f^i(N_v)| \leq \sum_{0 \leq i \leq h} \left(\left(\frac{4}{5}\right)^i |N_v| + \sum_{0 \leq j \leq i} \left(\frac{4}{5}\right)^j 4 \right) \\ &\leq |N_v| \left(\frac{1}{1 - \frac{4}{5}} \right) + 4 \sum_{0 \leq i \leq h} \left(\frac{1}{1 - \frac{4}{5}} \right) \leq 5|N_v| + 20h. \end{aligned}$$

Since the height h of the tree (N, R) corresponds to the number of bits needed to write the coordinates of the nodes in N , $h \in \mathcal{O}(\log n)$ and thus $|N| \in \mathcal{O}(n)$. ■

Note that the linearity constant obtained here is very large. Indeed, our goal here is to prove the linearity of the global algorithm, and not to provide a tight bound. With a more detailed analysis, the bound $|N| \leq 3|N_v| + 6h$ can be obtained for the number of nodes [13].

Theorem 1. *Given a word $w \in \Sigma^n$, determining if the path coded by w intersects itself is decidable in $\mathcal{O}(n)$.*

Proof. Lemma 3 implies that if the path w starts at $(0, 0)$ and stays in the first quadrant, then determining whether w intersects itself or no is decidable in linear time. All that is needed to adapt this solution to any word $w \in \Sigma^*$ is to use four graphs $\mathcal{G}_I, \mathcal{G}_{II}, \mathcal{G}_{III}, \mathcal{G}_{IV}$ simultaneously.



The cohesion between these graphs is ensured by *special nodes*, those representing points on axes. Since a point on an axis, distinct from the origin, is in exactly two quadrants, an additional link between the two nodes representing this point is added. These two nodes are *equivalent* since they represent the same point. The axis on which this point is located must also be stored in order to identify the quadrant in which the path enters. Switching from one graph to another is achieved with the following rules:

$$(II) a \leftrightarrow \bar{a}; \quad (III) a \leftrightarrow \bar{a}, b \leftrightarrow \bar{b}; \quad (IV) b \leftrightarrow \bar{b}.$$

This allows the processing of all coordinates as positive integers since their sign is determined by the quadrant. Consequently, each time a special node is created, its equivalent one is also created in the appropriate graph using the link between their fathers. ■

Performance issues and comparison. Among the many ways of solving the intersection problem, the naive sparse matrix representation that requires an $O(n^2)$ space and initialization time is eliminated in the first round. When efficiency is concerned, there are two well-known approaches for solving it: one may store the coordinates of the visited points and sort them, then check if two consecutive sets of coordinates are equal or not (we call this *sorting algorithm*). One may also store the sets of coordinates in an AVL-tree and check for each new set of coordinates if it is already present or not (the *AVL algorithm*). Let us first assume that the path $w \in \Sigma^n$ is not self intersecting. Then the length of the largest coordinate is $O(\log n)$. But the largest coordinate is also $\Omega(\log n)$ because if the path is not self intersecting, the minimum coordinates are obtained

when the points remain in a square centered on $(0, 0)$ with \sqrt{n} side length. Since $\log(\sqrt{n}) = \frac{1}{2} \log n$ the largest coordinate is also $\Omega(\log n)$. Thus the storage of the largest coordinate is in $\Theta(\log n)$ and the whole storage costs $\Theta(n \log n)$.

Sorting n ordered pairs can be done in $\Theta(n \log n)$ swaps or comparisons. But each swap or comparison costs $\Theta(\log n)$ clock ticks. Then the whole computation time is $\Theta(n \log^2(n))$. In our algorithm, the storage cost and computation time are both $\Theta(k)$ where k is the index of the second occurrence of the point appearing at least twice in the path or the path length if it is not self intersecting. Unlike the sorting algorithm, there is no need to store the whole path: the computation is performed dynamically. With our algorithm, storing the necessary data costs, both on average and in worst case, $\mathcal{O}(k)$ if k can be stored in a machine word and $\mathcal{O}(k \log k)$ otherwise. Similarly for the time complexity. We summarize :



Algorithm	Unified Cost RAM model		General Case	
	Time	Space	Time	Space
Sorting	$n \log n$	n	$n \log^2 n$	$n \log n$
AVL tree	$k \log k$	k	$k \log^2 k$	$k \log k$
Our	k	k	$k \log k$	$k \log k$

Consider the simpler problem of checking if a path is closed, that is if for each $\epsilon \in \Sigma$ we have $|w|_\epsilon = |w|_{\bar{\epsilon}}$. The cost of storing the number $|w|_\epsilon$ of occurrences of each elementary step is in $\mathcal{O}(1)$ if each of these numbers can be stored in a single machine word, or in $\mathcal{O}(\log(n_\epsilon))$ otherwise. Increasing or decreasing the number $|w|_\epsilon$ by 1 costs $\mathcal{O}(1)$ on average in both cases, and at worst $\mathcal{O}(1)$ for the first case and $\mathcal{O}(\log n)$ otherwise. The total time is hence $\mathcal{O}(n)$ in the first case and $\mathcal{O}(n \log n)$ otherwise.

There are other ways to deal with paths having negative coordinates. Indeed, since the property of being self intersecting or not is invariant by translation, it suffices to translate the path conveniently. This can be achieved by making one pass on the word w to determine the starting node \textcircled{s} as follows:

- (a) $\textcircled{s} \leftarrow (n, n)$, where $n = |w|$;
- (b) $\textcircled{s} \leftarrow (x, y)$ where x and y are determined from the extremal values.

In both cases it takes $\mathcal{O}(n)$ for reading the word, $\mathcal{O}(\log n)$ time and space to

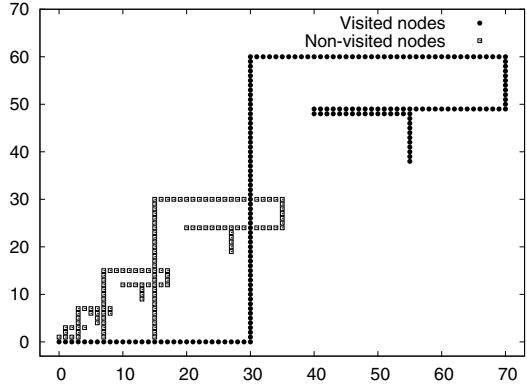


represent \textcircled{S} . Then the path is encoded in the radix-tree starting from \textcircled{S} . In our solution, we avoid the linear preprocessing for determining these values.

Numerical results. Our algorithms were implemented in C++ and tested on numerous examples. The results achieved for instance with $w_n = a^n b^n$ reveal a smaller linearity constant than the constant provided in the proof of Lemma 3, and confirmed their efficiency.

The radix-tree built for each word corresponds, as shown in Fig. 2, to points in the discrete plane $\mathbb{N} \times \mathbb{N}$. We provide on the right an illustration of the nodes involved in the radix-tree in the case of the word

$$w = a^{30}b^{60}a^{40}b^{11}a^{30}ba^{15}b^{10}.$$



5 Concluding Remarks

The first advantage of our algorithm is that ordering of edges can be used for avoiding labeling of both nodes and edges. Moreover, the neighboring relation T as presented is not implemented in its symmetric form. It could be easily done since each time a neighboring link $\textcircled{C} \xrightarrow{-\epsilon} \textcircled{Z}$ is added at line 9 of Algorithm 2, we can add its symmetric link $\textcircled{Z} \xrightarrow{-\epsilon} \textcircled{C}$ at constant cost. This does not change the overall complexity, and further analysis is required for determining if it is worthwhile. On the other hand, our algorithm is useful for solving a series of related problems in discrete geometry.

Determining if a path w crosses itself. When a node is visited twice, deciding whether the path crosses itself or not amounts to check local conditions, describing all the possible configurations (See [8] Section 4.1).

Determining if $w \in \Sigma^n$ is the Freeman chain code of a discrete figure. It suffices to check that the last visited node is the starting one. This does not penalize the linear algorithms for determining, for instance, if a discrete figure is digitally convex [7], or if it tiles the plane by translation [8]. In the case of a self intersecting path, it also allows the decomposition of a discrete figure in elementary components, not necessarily disjoint.

Node multiplicity. By replacing the “visited/unvisited” labeling of nodes with a counter (set to 0 when a node is created), the number of times a node is visited is computed by replacing the lines 4, 5, 6 and 7 in Algorithm 1 by the incrementation of this counter. Then, the obsolete line 10 must be removed.

Intersection of distinct paths. Given two distinct paths u and v of length bounded by $n = \max\{|u|, |v|\}$, with starting nodes \textcircled{S} and \textcircled{r} respectively, their

intersection is computed by constructing first the graph \mathcal{G}_u , inserting the node \mathcal{P} in \mathcal{G}_u , and constructing $\mathcal{G}_{u,v} = \mathcal{G}_u + \mathcal{G}_v$. Again the overall algorithm remains in $\mathcal{O}(n)$. As a byproduct of this construction, given two nonintersecting closed paths u and v , it is decidable whether the interior of u is included in the interior of v ; and consequently one may compute the exterior envelope of discrete figures.

Paths in higher dimension. The graph construction extends naturally to arbitrary d -tuples in $\mathbb{B}^* \times \dots \times \mathbb{B}^*$, for representing numbers in \mathbb{N}^d . Therefore, all the problems cited above can be treated in a similar way, by processing sequentially words on an alphabet $\Sigma_d = \{\epsilon_1, \bar{\epsilon}_1, \dots, \epsilon_d, \bar{\epsilon}_d\}$, of size $2d$. In the multidimensional case the trees used are no longer quadtrees but higher order trees, and in particular *octrees* for the 3-dimensional case.

All of these problems can be solved in linear time and space complexity.

Acknowledgements. The authors are grateful to the anonymous referee of our paper submitted to DGCi held in Szeged, who also reviewed the extended version appearing in [8], for bringing this problem to our attention. A preliminary version (in French) of the results presented here appears in the doctoral thesis of Xavier Provençal [9] who is supported by a scholarship from FQRNT (Québec).

References

1. Freeman, H.: On the encoding of arbitrary geometric configurations. IRE Trans. Electronic Computer 10, 260–268 (1961)
2. Freeman, H.: Boundary encoding and processing. In: Lipkin, B., Rosenfeld, A. (eds.) Picture Processing and Psychopictorics, pp. 241–266. Academic Press, New York (1970)
3. Brlek, S., Labelle, G., Lacasse, A.: A note on a result of Daurat and Nivat. In: De Felice, C., Restivo, A. (eds.) DLT 2005. LNCS, vol. 3572, pp. 189–198. Springer, Heidelberg (2005)
4. Brlek, S., Labelle, G., Lacasse, A.: Properties of the contour path of discrete sets. Int. J. Found. Comput. Sci. 17(3), 543–556 (2006)
5. Debled-Rennesson, I., Rémy, J.L., Rouyer-Degli, J.: Detection of the discrete convexity of polyominoes. Discrete Appl. Math. 125(1), 115–133 (2003)
6. Brlek, S., Lachaud, J.O., Provençal, X.: Combinatorial view of digital convexity. In: Coeurjolly, D., Sivignon, I., Tougne, L., Dupont, F. (eds.) DGCi 2008. LNCS, vol. 4992, pp. 57–68. Springer, Heidelberg (2008)
7. Brlek, S., Lachaud, J.O., Provençal, X., Reutenauer, C.: Lyndon+Christoffel = digitally convex. Pattern Recognition 42, 2239–2246 (2009)
8. Brlek, S., Provençal, X., Fédou, J.M.: On the tiling by translation problem. Discr. Appl. Math. 157, 464–475 (2009)
9. Provençal, X.: Combinatoire des mots, géométrie discrète et pavages. PhD thesis, D1715, Université du Québec à Montréal (2008)
10. Finkel, R., Bentley, J.: Quad trees: A data structure for retrieval on composite keys. Acta Informatica 4(1), 1–9 (1974)
11. Knuth, D.E.: The Art of Computer Programming, Sorting and Searching, vol. 3. Addison-Wesley, Reading (1998)
12. Lothaire, M.: Applied Combinatorics on Words. Cambridge University Press, Cambridge (2005)
13. Labbé, S.: Personal communication (2009)