# Towards a provably resilient scheme for graph-based watermarking 

Lucila M. S. Bento ${ }^{\text {a,b }}$, Davidson Boccardo ${ }^{\text {b }}$, Raphael C. S. Machado ${ }^{\text {b }}$, Vinícius G. Pereira de Sáa ${ }^{\text {a }}$, Jayme Luiz Szwarcfitera, ${ }^{\text {a,b }}$<br>${ }^{a}$ Instituto de Matemática, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brasil<br>${ }^{b}$ Instituto Nacional de Metrologia, Qualidade e Tecnologia, Rio de Janeiro, Brasil<br>${ }^{c}$ COPPE Sistemas, Universidade Federal do Rio de Janeiro, Rio de Janeiro, Brasil


#### Abstract

Digital watermarks have been considered a promising way to fight software piracy. Graph-based watermarking schemes encode authorship/ownership data as control-flow graph of dummy code. In 2012, Chroni and Nikolopoulos developed an ingenious such scheme which was claimed to withstand attacks in the form of a single edge removal. We extend the work of those authors in various aspects. First, we give a formal characterization of the class of graphs generated by their encoding function. Then, we formulate a lineartime algorithm which recovers from ill-intentioned removals of $k \leq 2$ edges, therefore proving their claim. Furthermore, we provide a simpler decoding function and an algorithm to restore watermarks with an arbitrary number of missing edges whenever at all possible. By disclosing and improving upon the resilience of Chroni and Nikolopoulos's watermark, our results reinforce the interest in regarding it as a possible solution to numerous applications.


Keywords: digital watermarking, permutation graphs, software security, robust algorithms, linear-time algorithms

[^0]
## 1. Introduction

The illegal reproduction of software has become a major concern for the industry. According to the Business Software Alliance, the commercial value of unlicensed software put into the world market in 2011 totaled 63.4 billion dollars [2]. To counter such practice, many promising methods have been devised, among which the idea of software watermarking.

The use of paper watermarks to prevent counterfeiting dates back to the thirteenth century. Generally speaking, watermarks are unique identifiers embedded into proprietary objects to enforce authenticity. In a digital object, particularly in a piece of software, a watermark may act not only as a certificate of authorship, but also as a means of tracing the original owner of the object, therefore discouraging piracy.

The first software watermark was proposed in 1996 by Davidson and Myrhvold [13], while the first watermarking scheme to exploit concepts of Graph Theory was formulated by Venkatesan, Vazirani and Sinha [22] in 2001. Their technique, whereby an integer was encoded as a special digraph disguised into the software's control-flow graph, was later patented [23]. Other original ideas, improvements and surveys on the available methods have been contributed by many authors ever since. See, for example, [9, 10, 11, 14, 20 , 25, 26].

Willing to prevent the timely retrieval of the encoded identification data, malicious agents may attempt to tamper with the watermark. A watermark solution is therefore only as secure as it is able to resist attacks of various sorts. Naturally, a lot of research has been put up lately towards developing more resilient solutions as well as strengthening existing ones. This paper pursues this latter goal.

We consider the graph-based watermarking scheme introduced by Collberg, Kobourov, Carter and Thomborson [9], and afterwards developed and improved upon by Chroni and Nikolopoulos in a series of papers [3, 4, $, 5,6,6,7]$. These latter authors proposed a watermark graph belonging to a subclass of the reducible permutation graphs introduced by the former authors. Though the mechanics of encoding and decoding the proposed watermark is well described in [5], such special subclass of reducible permutation graphs has not been fully characterized. Moreover, not much was known thus far about the resilience of Chroni and Nikolopoulos's graphs to malicious attacks, even though their ability to withstand single edge removals has been suggested without proof.

This paper is organized as follows. In Section 2, we present some preliminary concepts related to graph-based software watermarking, including the most common forms of attacks. In Section 3, we recall the watermark from Chroni and Nikolopoulos, and we state a number of structural properties, the proofs of which we delay until Section 7 for the sake of readability. In Section 4, we define and characterize the family of canonical reducible permutation graphs, which correspond to the watermarks produced by Chroni and Nikolopoulos's encoding function. In Section 5, we formulate linear-time algorithms to reconstruct the original digraph and recover the encoded data even if two edges are missing. The proof of one of the central results in that section, namely Theorem [24, is somewhat involved, and we dedicate a whole section to it towards the end of the paper. In Section 6, we propose a robust polynomial-time algorithm that, given a watermark with whatever number $k$ of missing edges, either recovers the encoded data or proves that the the watermark has become irremediably damaged. Finally, Sections 7 and 8 contain the postponed proofs for the properties stated in Section 3 and for Theorem [24, respectively. Section 9 concludes the paper with our final remarks.

Throughout the text, we let $V(G)$ and $E(G)$ respectively denote, as usual, the vertex set and edge set of a given graph $G$. Also, we let $N_{G}^{+}(v)$ and $N_{G}^{-}(v)$ be the sets of out-neighbors and in-neighbors of vertex $v$ in $G$, with $d_{G}^{+}(v)$ and $d_{G}^{-}(v)$ their respective sizes. If $J$ is a subset of either $V(G)$ or $E(G)$, then $G-J$ corresponds to the graph obtained from $G$ by the removal of $J$.

## 2. Graph-based software watermarking

Software watermarking schemes provide the necessary means of embedding identification data-typically a copyright notice or a customer numberinto a piece of software. We refer to the identification data as the identifier, and we may regard it as an integer, for simplicity. Watermarks are appropriate encodings of identifiers, and they can be broadly divided into two categories: static and dynamic [8]. The former are embedded in the code, whereas the latter are embedded into a program's execution state at runtime.

A static, graph-based watermarking scheme usually consists of four algorithms:

- an encoder, which converts the identifier into a graph - the watermark;
- a decoder, which extracts the identifier from the watermark;
- an embedder, a function whose input parameters are the software itself (either the binary code or the source code in some programming language), the intended watermark and possibly some secret key, and whose output is a modified software containing the watermark; and
- an extractor, which retrieves the watermark graph from the watermarked software.

To every computer program one can associate a directed graph representing the possible sequences of instructions (or, more precisely, of jump-free instruction blocks) during its execution. Such graph, called the control-flow graph (CFG) of the software [1], can be obtained by means of static analysis [17, 24]. What the embedder does is basically to insert dummy code into the program so that the intended watermark graph shows up as an induced subgraph of the CFG. The position of the watermark graph within the CFG is often determined as a function of a secret key. Knowledgeable of the secret key, the extractor retrieves that subgraph, which is then passed along to the decoding algorithm. Among the existing tools for embedding/extracting graph-based watermarks, we cite Collberg's SandMark project [12]. In this paper, we focus on the encoding/decoding algorithms described by Chroni and Nikolopoulos in [5].

Attacks. Among the several different kinds of attacks against graph-based watermarks, we list the following:

- additive attacks, in which other watermarks are inserted into the same object, generating ambiguity;
- subtractive attacks, in which the watermark is removed altogether; and
- distortive attacks, in which the watermark is modified to confound the decoder.

Additive and subtractive attacks can be precluded to a great extent by techniques of cryptography and software diversity [19]. On the other hand, distortive attacks-also known as jamming attacks-are more difficult to deal with and are arguably the most important attack model to be concerned about [22]. In some cases, the distortive attacker may even be able to reverse engineer the entire code and apply semantics-preserving modifications which modify the CFG structurally without affecting the software's functionalities.

## 3. The watermark by Chroni and Nikolopoulos

We recall the encoding algorithm described in [5]. The index of the first element in all considered sequences is 1 .

Let $\omega$ be a positive integer identifier, and $n$ the size of the binary representation $B$ of $\omega$. Let also $n_{0}$ and $n_{1}$ be the number of 0 's and 1's, respectively, in $B$, and let $f_{0}$ be the index of the leftmost 0 in $B$. The extended binary $B^{*}$ is obtained by concatenating $n$ digits 1 , followed by the one's complement of $B$ and by a single digit 0 . We let $n^{*}=2 n+1$ denote the size of $B^{*}$, and we define $Z_{0}=\left(z_{i}^{0}\right), i=1, \ldots, n_{1}+1$, as the ascending sequence of indexes of 0 's in $B^{*}$, and $Z_{1}=\left(z_{i}^{1}\right), i=1, \ldots, n+n_{0}$, as the ascending sequence of indexes of 1 's in $B^{*}$.

Let $S$ be a sequence of integers. We denote by $S^{R}$ the sequence formed by the elements of $S$ in backward order. If $S=\left(s_{i}\right)$, for $i=1, \ldots, t$, and there is an integer $k \leq t$ such that the subsequence consisting of the elements of $S$ with indexes less than or equal to $k$ is ascending, and the subsequence consisting of the elements of $S$ with indexes greater than or equal to $k$ is descending, then we say $S$ is bitonic. If all $t$ elements of a sequence $S$ are distinct and belong to $\{1, \ldots, t\}$, then $S$ is a permutation. If $S$ is a permutation of size $t$, and, for all $1 \leq i \leq t$, the equality $i=s_{s_{i}}$ holds, then we say $S$ is self-inverting. In this case, the unordered pair $\left(i, s_{i}\right)$ is called a 2 -cycle of $S$, if $i \neq s_{i}$, and a 1 -cycle of $S$, if $i=s_{i}$. If $S_{1}, S_{2}$ are sequences (respectively, paths in a graph), we denote by $S_{1} \| S_{2}$ the sequence (respectively, path) formed by the elements of $S_{1}$ followed by the elements of $S_{2}$.

Back to Chroni and Nikolopoulos's algorithm, we define $P_{b}=\left(b_{i}\right)$, with $i=1, \ldots, n^{*}$, as the bitonic permutation $Z_{0} \| Z_{1}^{R}$. Finally, the self-inverting permutation $P_{s}=\left(s_{i}\right)$ is obtained from $P_{b}$ as follows: for $i=1, \ldots, n^{*}$, element $s_{b_{i}}$ is assigned value $b_{n^{*}-i+1}$, and element $s_{b_{n^{*}-i+1}}$ is assigned value $b_{i}$. In other words, the 2-cycles of $P_{s}$ correspond to the $n$ unordered pairs of distinct elements of $P_{b}$ that share the same minimum distance to one of the extremes of $P_{b}$, that is, the pairs $(p, q)=\left(b_{i}, b_{n^{*}-i+1}\right)$, for $i=1, \ldots, n$. Since the central index $i=n+1$ of $P_{b}$ is the solution of equation $n^{*}-i+1=i$, element $b_{n+1}$ - and no other - will constitute a 1-cycle in $P_{s}$. We refer to such element of $P_{s}$ as its fixed element, and we let $f$ denote it.

The watermark generated by Chroni and Nikolopoulos's encoding algorithm [5] is a directed graph $G$ whose vertex set is $\{0,1, \ldots, 2 n+2\}$, and whose edge set contains $4 n+3$ edges, to wit: a path edge $(u, u-1)$ for $u=1, \ldots, 2 n+2$, constituting a Hamiltonian path that will be unique in $G$,


Figure 1: Watermark for identifier $\omega=43$.
and a tree edge from $u$ to $q(u)$, for $u=1, \ldots, n^{*}$, where $q(u)$ is defined as the vertex $v>u$ with the greatest index in $P_{s}$ to the left of $u$, if such $v$ exists, or $2 n+2$ otherwise. The rationale behind the name tree edge is the fact that such edges induce a spanning tree of $G \backslash\{0\}$.

Let us glance at an example. For $\omega=43$, we have $B=101011, n=6$, $n_{0}=2, n_{1}=4, f_{0}=2, B^{*}=1111110101000, n^{*}=13, Z_{0}=(7,9,11,12,13)$, $Z_{1}=(1,2,3,4,5,6,8,10), P_{b}=(7,9,11,12,13,10,8,6,5,4,3,2,1), P_{s}=(7,9$, $11,12,13,10,1,8,2,6,3,4,5)$ and $f=8$. The watermark graph associated to $\omega$ presents, along with the path edges in the Hamiltonian path $14,13, \ldots, 0$, the tree edges $(1,10),(2,8),(3,6),(4,6),(5,6),(6,8),(7,14),(8,10),(9,14)$, $(10,13),(11,14),(12,14)$ and $(13,14)$, as illustrated in Figure 1 .

In Section 4, we give a formal characterization of the class of such graphs, allowing for linear-time recognition and the formulation of efficient recovering algorithms against distortive attacks.

### 3.1. Structural properties

We now state a number of properties concerning the watermark from Chroni and Nikolopoulos and the special permutations they are associated to. These properties, whose proofs are given in Section 7 , set the basis for the characterization of the class of canonical reducible permutation graphs, which is given in Section 4, and for the recovering procedures described in Section 5 ,

For all properties stated below, let $G$ be the watermark graph associated to an identifier $\omega$ of size $n$, and let $P_{b}$ and $P_{s}$ be, respectively, the bitonic and the self-inverting permutations dealt with during the construction of $G$.

Property 1. For $1 \leq i \leq n$, the element $b_{n+i+1}$ in $P_{b}$ is equal to $n-i+1$, that is, the $n$ rightmost elements in $P_{b}$, from right to left, are $1, \ldots, n$.

Property 2. The elements whose indexes are $1, \ldots, n$ in $P_{s}$ are all greater than $n$.

Property 3. The fixed element $f$ satisfies $f=n+f_{0}$, unless the identifier $\omega$ is equal to $2^{k}-1$ for some integer $k$, whereupon $f=n^{*}=2 n+1$.

Property 4. In self-inverting permutation $P_{s}$, elements indexed $1, \ldots$, $f-n-1$ are respectively equal to $n+1, n+2, \ldots, f-1$, and elements indexed $n+1, n+2, \ldots, f-1$ are respectively equal to $1, \ldots, f-n-1$.

Property 5. The first element in $P_{s}$ is $s_{1}=n+1$, and the central element in $P_{s}$ is $s_{n+1}=1$.

Property 6. If $f \neq n^{*}$, then the index of element $n^{*}$ in $P_{s}$ is equal to $n_{1}+1$, and vice-versa. If $f=n^{*}$, then the index of element $n^{*}$ in $P_{s}$ is also $n^{*}$.

Property 7. The subsequence of $P_{s}$ consisting of elements indexed $1, \ldots$, $n+1$ is bitonic.

Property 8. For $u \leq 2 n,(u, 2 n+2)$ is a tree edge of watermark $G$ if, and only if, $u-n$ is the index of a digit 1 in the binary representation $B$ of the identifier $\omega$ represented by $G$.

Property 9. If $(u, k)$ is a tree edge of watermark $G$, with $k \neq 2 n+2$, then
(i) element $k$ precedes $u$ in $P_{s}$; and
(ii) if $v$ is located somewhere between $k$ and $u$ in $P_{s}$, then $v<u$.

## 4. Canonical reducible permutation graphs

This section is devoted to the characterization of the class of canonical reducible permutation graphs. After describing some terminology and proving some preliminary results, we define the class using purely graph-theoretical predicates. Then, we show it corresponds exactly to the set of watermarks produced by Chroni and Nikolopoulos's encoding algorithm [5]. Finally, we characterize it in a way that suits the design of a decoding algorithm which is simpler-for untampered with watermarks-and able to recover from removals of $k \leq 2$ edges in linear time.

A reducible flow graph [15, 16, 21] is a directed graph $G$ with source $s \in V(G)$, such that, for each cycle $C$ of $G$, every directed path from $s$ to $C$ reaches $C$ at the same vertex. It is well known that a reducible flow graph has at most one Hamiltonian cycle.

Definition 10. A self-labeling reducible flow graph is a directed graph $G$ such that
(i) $G$ presents exactly one directed Hamiltonian path $H$, hence there is a unique labeling function $\sigma: V(G) \rightarrow\{0,1, \ldots,|V(G)|-1\}$ of the vertices of $G$ such that the order of the labels along $H$ is precisely $|V(G)|,|V(G)|-1, \ldots, 0 ;$ and,
(ii) considering the labeling $\sigma$ as in the previous item, $N_{G}^{+}(0)=\emptyset$, $N_{G}^{-}(0)=\{1\}, N_{G}^{+}(|V(G)|-1)=\{|V(G)|-2\},\left|N_{G}^{-}(|V(G)|-1)\right| \geq 2$, and, for all $v \in V(G) \backslash\{0,|V(G)|-1\}, N_{G}^{+}(v)=\{v-1, w\}$, for some $w>v$.

From now on, without loss of generality, we shall take $\sigma$ for granted and assume the vertex set of any self-labeling reducible flow graph $G$ is the very set $V(G)=\{0,1, \ldots,|V(G)|-1\}$. By doing so, we may simply compare two vertices, e.g. $v>u$ (or $v$ greater than $u$, in full writing), whereas we would otherwise need to compare their images under $\sigma$, e.g. $\sigma(v)>\sigma(u)$.

Definition 11. The representative tree $T$ of a self-labeling reducible flow graph $G$ with Hamiltonian path $H$ has vertex set $V(T)=V(G) \backslash\{0\}$ and edge set $E(T)=E(G) \backslash E(H)$, where all edges are deprived of their orientation.

A representative tree $T$ is always regarded as a rooted tree whose root is $|V(G)|-1$, Moreover, it is regarded as an ordered tree, that is, for each $v \in V(T)$, the children of $v$ are always considered according to an ascending order of their labels. For $v \in T$, we denote by $N_{T}^{*}(v)$ the set of descendants of $v$ in $T$. Figure 2 depicts two representative trees.

Observation 12. The representative tree $T$ of a self-labeling reducible flow graph $G$ satisfies the max-heap property, that is, if vertex $u$ is a child of vertex $v$ in $T$, then $v>u$.

Proof: Direct from the way $T$ is rooted and from property (ii) in the definition of self-labeling reducible flow graphs, whereby the in-neighbors of $w$ in $G \backslash E(H)$ comprise only vertices $v<w$. We convey the idea that a representative tree $T$ satisfies the max-heap property by saying that $T$ is a descending, ordered, rooted tree.


Figure 2: Representative trees of the watermark graphs produced by Chroni and Nikolopoulos's encoding algorithm for identifiers (a) $\omega=31$ and (b) $\omega=43$ (the full watermark for $\omega=43$ is shown in Figure 11). It is easy to check that such graphs are self-labeling reducible flow graphs.

Definition 13. Let $S=\left(s_{i}\right), i=1, \ldots, 2 n+1$, be a self-inverting permutation. We say $S$ is canonical if:
(i) there is exactly one 1-cycle in $S$;
(ii) each 2-cycle $\left(s_{i}, s_{j}\right)$ of $S$ satisfies $1 \leq i \leq n$, for $s_{i}>s_{j}$;
(iii) $s_{1}, \ldots, s_{n+1}$ is a bitonic subsequence of $S$ starting at $s_{1}=n+1$ and ending at $s_{n+1}=1$.

Lemma 14. In any canonical self-inverting permutation, the fixed element $f$ satisfies $f \in[n+2,2 n+1]$.

Proof: By property (ii) of canonical self-inverting permutations, each 2-cycle of $S$ must contain at least one element whose index $i$ satisfies $1 \leq i \leq n$. From property (i), and given the size of $S$, it follows that the number of 2-cycles in $S$ is $n$, hence, by the pigeonhole principle, each and every 2 -cycle in $S$ contains exactly one such element $s_{i}$ with $1 \leq i \leq n$. But this means the other element in each 2 -cycle, namely $s_{j}$, satisfies $s_{j} \in[n+1,2 n+1]$. Since there are $n+1$ values in that range and only $n$ such elements $s_{j}$, there must be exactly one element $s_{k} \in[n+1,2 n+1]$ which is not part of a 2 -cycle, and therefore $s_{k}=f$. Now, by property (iii), $n+1=s_{1}$, hence $f \neq n+1$, and the lemma follows.

Let $T$ be a representative tree. The preorder traversal $P$ of $T$ is a sequence of its vertices that is recursively defined as follows. If $T$ is empty, $P$ is also empty. Otherwise, $P$ starts at the root $r$ of $T$, followed by the preorder traversal of the subtree whose root is the smallest child of $r$, followed by the preorder traversal of the subtree whose root is the second smallest child of $r$, and so on. The last (rightmost) element of $P$ is also referred to as the rightmost element of $T$.

Lemma 15. The preorder traversal of a representative tree $T$ is unique. Conversely, a representative tree $T$ is uniquely determined by its preorder traversal.

Proof: We use induction on $|V(T)|$. If $|V(T)| \leq 1$, the lemma holds trivially. Let $|V(T)|>1$, and let $v_{k}$ be the uniquely defined leaf of $T$ for which the path $v_{1}, \ldots, v_{k}$ from the root $v_{1}$ of $T$ to $v_{k}$ has the property that each $v_{i}$, for $1<i \leq k$, is the greatest vertex among the children of $v-1$. By the induction hypothesis, the preorder traversal $P^{\prime}$ of $T \backslash\left\{v_{k}\right\}$ is unique. Because $v_{k}$ is necessarily the rightmost vertex of $T$, the preorder traversal $P$ of $T$ is uniquely determined as $P^{\prime} \| v_{k}$.

Conversely, let $P$ be a preorder traversal of some representative tree $T$. If $|P| \leq 1$, there is nothing to prove. Otherwise, suppose the lemma holds for preorder traversals of size $\leq k$, and consider $|P|=k$. Let $v_{k}$ be the rightmost element of $P$. Clearly, $v_{k}$ must be a leaf of $T$, and also the rightmost (i.e., greatest) vertex among the children of its parent. Now define $P^{\prime}=$ $P-\left\{v_{k}\right\}$. By the induction hypothesis, there is a unique tree $T^{\prime}$ whose preorder traversal is $P^{\prime}$. Let $v_{k-1}$ be the rightmost element of $P^{\prime}$. We obtain $T$ from $T^{\prime}$, by making $v_{k}$ the rightmost child of the smallest ancestor $v_{j}$ of $v_{k-1}$ satisfying $v_{j}>v_{k}$, so $P$ is clearly the preorder traversal of $T$. Since no other parent for $v_{k}$ would be possible without breaking the ascending order of siblings in a representative tree, $T$ is uniquely defined by $P$.

The first element of the preorder traversal $P$ of a tree $T$ is always its root. If we remove the first element of $P$, the remaining sequence is said to be the root-free preorder traversal of $T$.

We can now define the class of canonical reducible permutation graphs.
Definition 16. A canonical reducible permutation graph $G$ is a self-labeling reducible flow graph on $2 n+3$ vertices, for some integer $\geq 1$, such that the root-free preorder traversal of the representative tree of $G$ is a canonical self-inverting permutation.

Lemma 17. If $G$ is a watermark instance produced by Chroni and Nikolopoulos's encoding algorithm [5], then $G$ is a canonical reducible permutation graph.

Proof: Recall, from Section 3, that the watermark graph $G$ associated to identifier $\omega$, whose binary representation $B$ has size $n$, is constructed with vertex set $V(G)=\{0, \ldots, 2 n+2\}$ and an edge set $E(G)$ which can be partitioned into path edges and tree edges in such a way that all conditions in the definition of self-labeling reducible flow graphs are satisfied, as can be easily checked. Now, by Property 9 of Chroni and Nikolopoulos's watermarks (see Section 3.1), the tree edges of $G$ constitute a representative tree $T$ of $G$ whose root-free preorder traversal is precisely the self-inverting permutation $P_{s}$ determined by the encoding algorithm from [5] as a function of $B$. Consequently, what is left to prove is that $P_{s}$ is canonical. The first condition to $P_{s}$ being canonical is asserted by Property 3 in Section 3.1(the fixed element $f$ corresponds to the unique 1-cycle in $P_{s}$ ); the second condition is given by Property 2; and, finally, Properties 5 and 7 fulfill the third condition, therefore $P_{s}$ is canonical.

Lemma 18. If $G$ is a canonical reducible permutation graph, then $G$ is the watermark produced by Chroni and Nikolopoulos's encoding algorithm [5] for some integer identifier $\omega$.

Proof: Let $G$ be a canonical reducible permutation graph, and $T$ its representative tree. By Lemma 15, $T$ is uniquely defined by its preorder traversal $P$. We show that $P$ corresponds to the self-inverting permutation $P_{s}$ generated by the encoding algorithm of [5] (please refer to Section 3 for details) when computing the watermark for some integer identifier $\omega$. By definition, $P=\left(s_{i}\right), i=1, \ldots, 2 n+1$, is a canonical self-inverting permutation presenting a single 1 -cycle $f$ and a number $n$ of 2 -cycles $(p, q)$. Those 2-cycles $(p, q)$ define exactly one bitonic permutation $P_{b}=\left(b_{j}\right), j=1, \ldots, 2 n+1$ satisfying Property 1 of Chroni and Nikolopoulos's watermarks with
(i) $b_{n+1}=f$, and,
(ii) for all $j \in\{1, \ldots, n\}, b_{j}=p$ if and only if $b_{2 n+1-j}=q$.

Such bitonic permutation $P_{b}$ can be regarded as $Z_{0} \| Z_{1}^{R}$ by assigning to $Z_{0}$ the prefix of $P_{b}$ comprising its maximal ascending subsequence, and now the
indexes of 0 's and 1's in the extended binary $B^{*}$ are totally determined. We proceed by extracting the binary $B$ that is the one's complement of the subsequence of $B^{*}$ with digits from the $(n+1)$ th to the $(2 n)$ th position in $B^{*}$. Regarding $B$ as the binary representation of a positive integer $\omega$, the image of such $\omega$ under the encoding function of Chroni and Nikolopoulos is isomorphic to $G$.

We proceed to the last definitions before we can give an appropriate, algorithmic-flavored characterization of canonical reducible permutation graphs. Let $T$ be the representative tree of some canonical reducible permutation graph $G$, and $P$ a canonical self-inverting permutation corresponding to the root-free preorder traversal of $T$. We refer to the fixed element $f$ of $P$ also as the fixed element (or vertex) of both $G$ and $T$. Similarly, the 2-cyclic elements of $P$ correspond to cyclic elements (or vertices) of both $G$ and $T$. A vertex $v \in V(T) \backslash\{2 n+2\}$ is considered large when $n<v \leq 2 n+1$; otherwise, $v \leq n$ and $v$ is dubbed as small. Denote by $X, Y$, respectively, the subsets of large and small vertices in $T$, so $|X|=n+1$ and $|Y|=n$. By Lemma 14, $f \in X$. We then define $X_{c}=X \backslash\{f\}=\left\{x_{1}, \ldots, x_{n}\right\}$ as the set of large cyclic vertices in $T$.

Definition 19. A representative tree $T$ is a Type-1 tree - see Figure 3(a) - when
(i) $n+1, n+2, \ldots, 2 n+1$ are children of the root $2 n+2$ in $T$; and
(ii) $1,2, \ldots, n$ are children of $2 n$.

Definition 20. A representative tree $T$ is a Type- 2 tree relative to $f-$ see Figure $3(b)$ - when
(i) $n+1=x_{1}<x_{2}<\ldots<x_{\ell}=2 n+1$ are the children of $2 n+2$, for some $\ell \in[2, n-1]$;
(ii) $x_{i}>x_{i+1}$ and $x_{i}$ is the parent of $x_{i+1}$, for all $i \in[\ell, n-1]$;
(iii) $1,2, \ldots, f-n-1$ are children of $x_{n}$;
(iv) $x_{i}=n+i$, for $1 \leq i \leq f-n-1$;
(v) $f$ is a child of $x_{q}$, for some $q \in[\ell, n]$ satisfying $x_{q+1}<f$ whenever $q<n$; and


Figure 3: (a) A Type-1 representative tree. (b) A Type-2 representative tree.
(vi) $N_{T}^{*}(f)=\{f-n, f-n+1, \ldots, n\}$ and $y_{i} \in N_{T}^{*}(f)$ has index $x_{y_{i}}-f+1$ in the preorder traversal of $N_{T}^{*}[f]$.

Lemma 21. If $y_{r}$ is the rightmost vertex of a Type-2 representative tree $T$ relative to some $f \neq 2 n+1$, then $y_{r}$ is equal to the number $\ell$ of children of the root $2 n+2$ in $T$.

Proof: By the definition of a Type-2 representative tree, the only nonleaf child of the root $2 n+2$ of $T$ is its rightmost child $x_{\ell}$, therefore each child $x_{i}$ of $2 n+2$, for $1 \leq i \leq \ell$, appears precisely at the $i$ th position in the root-free preorder traversal $P$ of $T$. Since, by definition, $P$ is self-inverting, and $y_{r}$ is the last, $(2 n+1)$ th element of $P$, it follows that $y_{r}$ must be equal to the index of $2 n+1=x_{\ell}$ in $P$, that is, $y_{r}=\ell$.

The following theorem characterizes canonical reducible permutation graphs in terms of the above defined trees. Such characterization is crucial for the remainder of the paper, since its straightforward conditions can be checked in linear time and give rise to the decoding and recovering algorithms that will come next.

Theorem 22. A digraph $G$ is a canonical reducible permutation graph if, and only if, $G$ is a self-labeling reducible flow graph and
(i) the fixed element of $G$ is $2 n+1$ and $G$ has a Type-1 representative tree; or
(ii) the fixed element of $G$ belongs to $[n+2,2 n]$ and $G$ has a Type-2 representative tree.

Proof: Let $G$ be a canonical reducible permutation graph and $T$ its representative tree. By definition, $G$ is a self-labeling reducible flow graph. Let $P=s_{1}, \ldots, s_{2 n+1}$ be the root-free preorder traversal of $T$, hence a canonical self-inverting permutation, also by definition. This means, among other things, that $P$ has a unique fixed element $f$, and that $P^{\prime}=s_{1}, \ldots, s_{n+1}$ is a bitonic subsequence of $P$. Since $T$ is descending, it follows that the prefix $A$ of $P^{\prime}$ constituting its maximal ascending subsequence must comprise solely vertices that are children of the root $2 n+2$, the rightmost of which certainly being $2 n+1$.

First, let $f=2 n+1$. Since $f$ constitutes a 1-cycle of $P, f$ must occupy the rightmost, $(2 n+1)$ th position in $P$, hence $f$ is a leaf of $T$. Furthermore, by Property 4 of canonical reducible permutation graphs, it follows that $P^{\prime}$ consists of elements $n+1, n+2, \ldots, 2 n, 1$, hence $A=n+1, n+2, \ldots, 2 n$, and these vertices are therefore children of $2 n+2$ in $T$. Now, again by Property 4, elements $1, \ldots, n$ appear, in this order, to the right of $A$ in $P$. Considering that $P$ is a preorder traversal and a representative tree satisfies the max-heap property, we conclude that vertices $1, \ldots, n$ can only be children of $2 n$, hence $T$ is a Type- 1 tree, as required.

Next, suppose $f<2 n+1$. By Lemma 14, it follows that $f \in[n+2,2 n]$. We already know that the children of $2 n+2$ are the vertices of $A$. Let $D$ be the subset formed by the remaining vertices of $P^{\prime}$. Clearly, the vertices of $D$ must appear in descending order. Since $T$ satisfies the max-heap property and $P$ is a preorder traversal of $T$, it follows that the largest vertex of $D$ is a child of $2 n+1$, and subsequently each vertex in $D$ is the parent in $T$ of the vertex placed to its left along the sequence $D$. Again, because $T$ satisfies the max-heap property, $f \in[2 n+2,2 n]$ must be the child of the smallest vertex $x_{q} \in D \cup\{2 n+1\}$ satisfying $x_{q}>f$. Let us again examine the ascending subsequence $A$. We know that the first vertex of $A$ is $n+1$. Suppose the leading $k$ vertices of $A$ are $n+1, n+2, \ldots, n_{k}$, for some $k$. Because $P$ is self-inverting, it follows that the vertices $1,2, \ldots, k$ must be the children of the last (i.e., smallest) vertex of $D$, and $k=f-n-1$ by Property 4. It remains solely to describe how the remaining small vertices, namely $f-n, f-n+1, \ldots, n$, are placed in $T$. Since they appear after $f$ in $P$, it can only be that this subset comprises exactly the descendants $N_{T}^{*}(f)$ of $f$ in $T$. Each of the vertices $y \in N_{T}^{*}(f)$ constitute a 2-cycle with some
vertex $x$ belonging to the bitonic subsequence $P^{\prime}$, hence the index of $y$ in $P$ is exactly $x$, and all the conditions for a Type- 2 tree have thus been verified.

Conversely, let $G$ be a self-labeling reducible flow graph. First, suppose that (i) applies and let $T$ be the corresponding Type-1 representative tree. Then the root-free preorder traversal $P$ of $T$ is

$$
n+1, n+2, \ldots, 2 n, 1,2, \ldots, n, 2 n+1
$$

Regarding $P$ as a permutation of $\{1, \ldots, 2 \mathrm{n}+1\}$, we observe that $2 n+1$ is the only fixed vertex on it; for $1 \leq i \leq n$, each element $n+i$ of $P$ has index $i$, while $i$ has index $n+i$, and $n+1, n+2, \ldots, 2 n, 1$ form a bitonic subsequence of $P$. Consequently, $P$ is a canonical self-inverting permutation, and $G$ is a canonical reducible permutation graph.

Finally, suppose (ii) applies. Let $T$ be the corresponding Type- 2 representative tree relative to some $f \in[n+2,2 n]$. The root-free preorder traversal $P$ of $T$ consists of

$$
x_{1}, \ldots, x_{\ell}, x_{\ell+1}, \ldots, x_{q}, x_{q+1}, \ldots, x_{n}, 1,2, \ldots, f-n-1, f, P\left(N_{T}^{*}(f)\right)
$$

where $x_{1}=n+1 ; x_{\ell}=2 n+1 ; x_{i}=n+i$ for $1 \leq i \leq f-n-1 ; x_{1}, x_{2}, \ldots, x_{n}, 1$ is a bitonic subsequence of $P$; and $P\left(N_{T}^{*}(f)\right)$ denotes the preorder traversal of the vertices of $N_{T}^{*}(f)$, in which each $y_{i} \in N_{T}^{*}(f)$ has index $x_{y_{i}}-f+1$. Observe that, for $1 \leq i \leq f-n-1,(n+i, i)$ constitutes a 2 -cycle in $P$. Moreover, for $f-n \leq i \leq n$, vertex $x_{i}$ forms a 2 -cycle with an element $y_{j} \in N_{T}^{*}(f)$. All conditions have been met, thus $P$ is a canonical self-inverting permutation and $G$ is a canonical reducible permutation graph.

Corollary 23. The recognition of canonical reducible permutation graphs can be achieved in linear time.

Proof: Direct from Theorem 22 and from the definitions of self-labeling reducible flow graphs, Type-1 and Type-2 representative trees, all of whose conditions can be verified in linear time easily.

## 5. Linear-time decoding ( $k \leq 2$ missing edges)

In this section, we analyze the effects of a distortive attack against a watermark (i.e., a canonical reducible permutation graph) $G$ from which

```
Algorithm 1 reconstruct_Hamiltonian_path \(\left(G^{\prime}\right)\)
input: a damaged watermark \(G^{\prime}\) with \(2 n+3\) vertices and two missing edges
output: the unique Hamiltonian path \(H\) in the original watermark \(G\)
```

1. Let $V_{0}$ be the set of all vertices with degree zero in $G^{\prime}$.
2. Let $\mathcal{H}$ be the set of all Hamiltonian path candidates obtained by a call to plug_next_subpath $\left(G^{\prime}, V_{0}, \emptyset\right)$

## 3. for each Hamiltonian path candidate $H \in \mathcal{H}$ do if validate_labels $\left(G^{\prime}, H\right)$ then <br> return $H$

$k \leq 2$ edges were removed. Note that the unique Hamiltonian path $H$ of $G$ may have been destroyed by the attack. The knowledge of $H$ is crucial for determining the labels of the vertices (they range from $2 n+2$ to 0 along $H$ ). Our first task is therefore to determine whether any path edges are missing from $G$, so we can restore $H$ and label the vertices accordingly.

### 5.1. Reconstructing the Hamiltonian path

The algorithm given in pseudocode as Algorithm retrieves the unique Hamiltonian path $H$ of a (possibly damaged) watermark $G^{\prime}$, that is, a graph isomorphic to a canonical reducible permutation graph $G$ minus $k \leq 2$ edges. It employs two subroutines presented separately: plug_next_subpath and validate_labels. The algorithm itself is straightforward. It basically builds Hamiltonian path candidates for $G$ (possibly by reinserting some edges) and tests whether the vertex labeling implied by each such candidate satisfies some conditions. It returns the one and only candidate which passes the test.

The procedure plug_next_subpath, given as Algorithm 2, takes as input a graph $G^{\prime}$, a path $Q$ with $V(Q) \cap V(G)=\emptyset$, and an output list $\mathcal{H}$, where (restored) Hamiltonian path candidates of $G^{\prime}$ will be placed after being concatenated to (the left of) a copy of $Q$. It starts by determining a set $S$ of subpath heads. This set comprises every vertex $s \in V\left(G^{\prime}\right)$ satisfying $d_{G^{\prime}}^{+}(s)=0$, in case $Q=\emptyset$, or $d_{G^{\prime}}^{+}(s) \leq 1$, otherwise. Then it computes the collection of all maximal backward-unbifurcated paths of $G^{\prime}$ reaching $S$ (or $S$-bups). An $S$-bup of $G^{\prime}$ is a path $v_{j}, v_{j-1}, \ldots, v_{1}$ such that

- $v_{1}=s$, for some $s \in S$;
- $\left(v_{k}, v_{k-1}\right) \in E\left(G^{\prime}\right)$, for $2 \leq k \leq j$; and
- the in-degree of $v_{k}$ in $G^{\prime}-\left\{v_{1}, \ldots, v_{k-1}\right\}$ satisfies

$$
d_{G^{\prime}-\left\{v_{1}, \ldots, v_{k-1}\right\}}^{-}\left(v_{k}\right)= \begin{cases}1, & \text { if } 1 \leq k \leq j-1 \\ 0, & \text { if } k=j\end{cases}
$$

In other words, starting from some subpath head $s=v_{1}$, the procedure builds a directed path $Q$ in backwards fashion by concatenating an in-neighbor of $v_{k}$ to the left of $v_{k}$, for $k \geq 1$, whenever the indegree of $v_{k}$ is 1 in the graph induced by all vertices which have not yet been incorporated to the path. It carries on iteratively this way until, for some $j$, the indegree of $v_{j}$ in the aforementioned graph is either zero, whereupon it adds the path so obtained to a list of $S$-bups, or greater than one, whereupon it discards the current path. The rationale behind it is that a backward bifurcation on $v_{j}$ means there are two vertices, say $u$ and $w$, which have not yet been added to the path, both of which are in-neighbors of $v_{j}$. Since at most one of them, say $u$, may be the tail of a path edge pointing to $v_{j}$ in $Q$, the other one, $w$, will be the tail of a tree edge pointing to $v_{j}$, which is not acceptable since $w$ will be to the left of $v_{j}$ in the path. Whichever the case, the algorithm starts anew with another subpath head $s \in S$ until all of them have been considered and the list of $S$-bups is fully populated. Finally, it appends each $S$-bup $Q^{\prime}$, one at each time, to the left of $Q$ (by adding a plausible path edge e $\notin E\left(G^{\prime}\right)$ from the rightmost vertex in $Q^{\prime}$ to the leftmost vertex in $Q$ ) and performs one of two possible actions:

- if $V\left(Q^{\prime}\right)=V\left(G^{\prime}\right)$ (i.e., if $Q^{\prime}$ is a Hamiltonian path of $G^{\prime}$ ), than it adds the new path $Q^{\prime} \| Q$ to the output list $\mathcal{H}$;
- otherwise, it makes a recursive call to plug_next_subpath with parameters $G^{\prime}-V\left(Q^{\prime}\right), Q^{\prime} \| Q$, and the output list $\mathcal{H}$.

When all $S$-bups have been considered, it returns $\mathcal{H}$.
If $H$ is a path, then we indicate the $j$ th element of $H$ (from right to left, starting at $j=0$ ) by $H[j]$.

The second subroutine invoked by Algorithm 1 is called validate_labels, shown in pseudocode as Algorithm 3. It takes as parameters a watermark

```
Algorithm 2 plug_next_subpath \(\left(G^{\prime}, Q, \mathcal{H}\right)\)
input: a graph \(G^{\prime}\), a path \(Q\) with \(V(Q) \cap V\left(G^{\prime}\right)=\emptyset\),
    and an output list \(\mathcal{H}\)
output: an updated \(\mathcal{H}\) containing all \(Q^{\prime}| | Q\)
    where \(Q\) is a Hamiltonian path of \(G^{\prime}\) (plus \(k \geq 0\) extra edges)
    ending at a vertex of degree \(d \leq 1\) (or \(d=0\), if \(Q\) is empty)
```

1. Let $S \leftarrow\left\{s \in V\left(G^{\prime}\right): d_{G^{\prime}}^{+}(s)=0\right\}$.
if $Q \neq \emptyset$ then $S \leftarrow S \cup\left\{s \in V(G): d_{G^{\prime}}^{+}(s)=1\right\}$
2. for each $s \in S$ do
```
    \(v \leftarrow s\)
    \(Q^{\prime} \leftarrow s\)
    while \(\left|N_{G^{\prime}-\left(V\left(Q^{\prime}\right)-\{v\}\right)}^{-}(v)\right|=1\) do
        \(v \leftarrow\) the unique element in \(N_{G^{\prime}-\left(V\left(Q^{\prime}\right)-\{v\}\right)}^{-}(v)\)
        \(Q^{\prime} \leftarrow v \| Q^{\prime}\)
    if \(\left|Q^{\prime}\right|=\left|V\left(G^{\prime}\right)\right|\) then
        \(\mathcal{H} \leftarrow \mathcal{H} \cup\left\{Q^{\prime}| | Q\right\}\)
    else if \(\left|N_{G^{\prime}-\left(V\left(Q^{\prime}\right)-\{v\}\right)}^{-}(v)\right|=0\) then
        plug_next_subpath \(\left(G^{\prime}-V\left(Q^{\prime}\right), Q^{\prime} \| Q, \mathcal{H}\right)\)
    else discard \(Q^{\prime} / /\) a backward bifurcation was found
```

$G^{\prime}$ (with two missing edges) and a candidate Hamiltonian path $H$. First, it determines the set $H^{*}=E(H) \backslash E\left(G^{\prime}\right)$ of the $k \leq 2$ plausible path edges that were required by $H$. It then checks whether it is possible to obtain a valid canonical reducible permutation graph $G$ through the insertion of $H^{*}$ and some set of $2-k$ tree edges into $G^{\prime}$. It does so by testing the following necessary conditions, where $T$ denotes the representative tree of $G$ :
(1) vertices $H[2 n+1]$ and $H[n+1]$ must be in-neighbors of $H[2 n+2]$ in $G$;
(2) the out-degree of vertices $H[2 n+1], \ldots, H[1]$ must be 2 in $G$, and the out-degree of $H[2 n+2]$ must be 1 ;
(3) the number of tree edges that would have to be inserted into $G$ so that the two previous conditions are met must not exceed $2-\left|H^{*}\right|$; and, finally,
(4) if vertices $H[1]$ and $H[n]$ are not siblings in $T$, then vertex $H$ [1] must be a child of the $n$th descendant of the root $H[2 n+2]$ which is a large vertex (i.e., the $n$th descendant of the root, counted right to left, among those whose indexes in $H$ are greater than or equal to $n+1$ ); moreover, the index in $H$ of the rightmost vertex in the preorder traversal of $T$ must correspond to the number of children of $2 n+2$ in $T$.

The first condition above appears in the definition of both Type-1 and Type-2 representative trees (Definitions (19) and 20), hence its necessity comes directly from Theorem [22. The second condition is due to Definition 16 and from the second property in the definition of self-labeling reducible flow graphs (Definition (10). The third condition obviously comes from the fact that 2 edges were removed from $G$, and $\left|H^{*}\right|$ edges have already been (re)inserted at this point. Finally, the last condition is due to property (iii) in the definition of Type-2 representative trees and to Lemma 21, It certainly applies to Type-2 representative trees only, which is precisely the case where vertices $H[1]$ and $H[n]$ are not siblings in the representative tree, by definition. We remark that, if $H$ is indeed the unique Hamiltonian path of a canonical reducible permutation graph $G$, then, for all $v \in V(G)$, the canonical label $v$ satisfies $v=H[v]$.

Theorem 24. Algorithm 1 correctly retrieves the original, unique Hamiltonian path from a canonical reducible permutation graph on $2 n+3$ vertices from which $k \leq 2$ edges were removed. It runs in $O(n)$ time.

The proof of Theorem 24 is somewhat involved and unfortunately demands some case analysis. We therefore postpone it until Section 8 .

### 5.2. Determining the fixed vertex

Suppose the watermark $G$ has been attacked, which resulted in a damaged watermark $G^{\prime}$, where two unknown edges are missing. Now we shall recognize the fixed vertex of the original watermark, given the damaged one. Getting to know the fixed vertex of $G$ will play a crucial role in retrieving the missing tree edges and consequently restoring the original identifier $w$ encoded by $G$.

We describe some characterizations that lead to an efficient computation of the fixed vertex $f$ of $G$. Let $T$ be the representative tree of the original watermark $G$. We consider the case where the two edges that have been removed belong to $T$. Denote by $F$ the forest obtained from $T$ by the removal of two edges. First, we consider the case $f=2 n+1$.

```
Algorithm 3 validate_labels \(\left(G^{\prime}, H\right)\)
input: a graph \(G^{\prime}\), with \(\left|V\left(G^{\prime}\right)\right|=2 n+3\),
    and a Hamiltonian path candidate \(H\), with \(\left|E(H) \backslash E\left(G^{\prime}\right)\right| \leq 2\)
output: True, if the labeling of \(V\left(G^{\prime}\right)\) implied by \(H\) is valid; False, otherwise
1. Label the vertices of \(G^{\prime}\) in such a way that \(H=2 n+2,2 n+1, \ldots, 0\).
Let \(H^{*} \leftarrow E(H) \backslash E\left(G^{\prime}\right)\), and insert \(H^{*}\) into \(G^{\prime}\) obtaining \(G^{\prime \prime}\).
Let also \(F\) be the forest obtained from \(G^{\prime \prime}\) by the removal of all (path)
edges in \(H\), as well as the isolated vertex 0 , and let missing_edges \(\leftarrow 0\).
2. for each \(v \in\{n+1,2 n+1\}\) do
\[
\text { if }(v, 2 n+2) \notin E\left(G^{\prime \prime}\right) \text { then }
\]
\[
\text { if } d_{G^{\prime \prime}}^{+}(v)=2 \text { then return False }
\]
\[
E\left(G^{\prime \prime}\right) \leftarrow E\left(G^{\prime \prime}\right) \cup\{(v, 2 n+2)\} ; \text { missing_edges }+=1
\]
for each \(v \in\{1, \ldots, 2 n+1\}\) do
if \(d_{G^{\prime \prime}}^{+}(v)<2\) then missing_edges \(+=1\)
if missing_edges \(>2-\left|H^{*}\right|\) then return False
```

3. if vertices 1 and $n$ are siblings in $F$ then

Let $r$ be the rightmost vertex in the preorder traversal of $F$.
if $r>d_{G^{\prime \prime}}^{-}(2 n+2)+$ missing_edges then return False
Let $x$ be the length of the unique path from 1 to $2 n+2$ in $F$. if $r+x-2 \neq n$ then return False
4. return True

Theorem 25. Let $F$ be a forest obtained from the representative tree $T$ by removing two edges, where $n>2$. Then $f=2 n+1$ if, and only if,

1. vertex $2 n+1$ is a leaf of $F$; and
2. the $n$ small vertices of $G^{\prime}$ are children of $2 n$ in $F$, with the possible exception of at most two of them, in which case they must be isolated vertices.

Proof: From Theorem [22, we know that, when $f=2 n+1, f$ is the rightmost vertex of $T$, hence a leaf of $F$, implying the necessity of condition (i). Again by Theorem 22, the small vertices of $T$ must immediately follow the rightmost


Figure 4: (a-c) Conditions (i), (ii) and (iii) of Theorem 26, respectively.
cyclic vertex of $T$, namely $2 n$. Since two edges have been deleted from $T$, it follows that all small vertices are children of $2 n$ in $F$, with the possible exception of at most two of them, which then became isolated vertices, so condition (ii) is also necessary.

Conversely, suppose conditions (i) and (ii) hold, and assume $f \neq 2 n+1$. Then the second case covered by Theorem 22 applies for $T$. If $f=2 n$, we know from $n>2$ that vertex $2 n+1$ has at least 3 children in $T$, making it impossible for $2 n+1$ to become a leaf of $F$ by the removal of only two edges, therefore contradicting condition (i). If, on the other hand, $f<2 n$, then, again by Theorem [22, vertex $2 n+1$ cannot have any small children, contradicting condition (ii). Therefore $f>2 n$, implying $f=2 n+1$.

Next, we characterize the case $f<2 n+1$. Figure 4 helps to visualize the three conditions of the theorem.

Theorem 26. Let $F$ be a forest obtained from the representative tree $T$ of watermark $G$ by removing two of its edges, and let $x \leq 2 n$ be a large vertex of $T$ which is not a child of $2 n+2$. Then $x$ is the fixed vertex $f$ of $G$ if, and only if,
(i) the large vertex $x$ has a sibling $z$ in $F$, and $x>z$; or
(ii) the subset of small vertices $Y^{\prime} \subset Y, Y^{\prime}=\{x-n, x-n+1, \ldots, n\}$ can be partitioned into at most two subsets $Y_{1}^{\prime}, Y_{2}^{\prime}$, such that $\emptyset \neq Y_{1}^{\prime}=N_{F}^{+}(x)$ and $Y_{2}^{\prime}$ is the vertex set of one of the trees which form $F$; or, whenever the previous conditions do not hold,
(iii) the large vertex $x$ is the rightmost vertex of one of the trees of $F$, while the rightmost vertices of the remaining trees are all small vertices.

Proof: For the sufficiency of condition (i), let $x$ be a large vertex of $F$, $z$ a sibling of $x$ in $F$ and $x_{q}$ their parent. By Theorem 22, the only large vertex of $T$ which is not a child of $2 n+2$ and has some sibling $z$ is precisely the fixed vertex $f$. Clearly, the removal of edges of $T$ cannot create new vertices having this property. Furthermore, $x_{q} \notin\left\{x_{n}, 2 n+1\right\}$ implies that $f$ has a unique sibling $x_{q+1}$, hence $f>x_{q+1}$ according to the ascending order of siblings in $F$, whereas $x_{q} \in\left\{x_{n}, 2 n+1\right\}$ implies every sibling $y$ of $f$ is a small vertex, hence $f>y$. Consequently, $x=f$.

Now suppose condition (ii) holds. First, assume that $Y_{2}^{\prime}=\emptyset$. In this case, $Y^{\prime}=Y_{1}^{\prime}=\{x-n, x-n+1, \ldots, n\}=N_{F}^{*}(x)$. Again, according to Theorem 22, we can locate a unique vertex $f$ fulfilling this property, implying $x=f$. In addition, when $Y_{2}^{\prime} \neq \emptyset$, we can again select a unique vertex $f$, where $N_{F}^{*}(f) \cup Y_{2}^{\prime}=Y^{\prime}$. Thus, $x=f$ indeed.

Finally, assume neither condition (i) nor condition (ii) hold. Because (i) is not satisfied, we have that either $x_{q} \notin\left\{x_{n}, 2 n+1\right\}$, and the edge from $x_{q}$ to one of its children has been deleted; or $x_{q} \in\left\{x_{n}, 2 n+1\right\}$, and the edge $\left(x_{q}, f\right)$ has been deleted. Additionally, since (ii) is not satisfied, $f$ must have a unique child $y$, and the edge $(f, y)$ has also been removed. Next, assume that, in such a context, condition (iii) is verified. For the sake of contradiction, suppose the theorem is false, so that $x \neq f$. Since $x \neq 2 n+1$ and $x$ is not a child of $2 n+2$, it follows that it must be a descending vertex, whereupon the fact that $x$ is the rightmost vertex of the tree of $F$ containing it implies that $x$ is a leaf of $F$. Now the latter implies that the edge $\left(x_{q}, x\right)$ of $T$ has been removed, where $x_{q}$ is the parent of $x$ in $T$. Because condition (i) is not satisfied, at least one edge has been removed from $T$, and because condition (ii) is not satisfied, at least one more edge has been deleted from $T$. Since no more than two edges overall have been removed, we conclude that the assumption is false, and therefore, here again, $x=f$.

Conversely, assume that $x$ is the large vertex of $F$ satisfying $x=f$. We prove that condition (i) or condition (ii) holds, otherwise condition (iii) is satisfied.

Let $x_{q}$ be the parent of $x=f$ in $T$. If $x_{q}$ has at least two children in $F$, then $f$ is larger than its siblings, by Theorem [22, and condition (i) holds. Alternatively, if $f$ is not a leaf of $F$, then the set $Y^{\prime}=\{x-n, x-n+1, \ldots, n\}$ either satisfies $Y^{\prime}=N_{F}^{*}(f)$ or it can be split into two subsets $Y_{1}^{\prime} \cup Y_{2}^{\prime}=Y^{\prime}$, where $Y_{1}^{\prime}=N_{F}^{*}(f)$ and $Y_{2}^{\prime}$ is the vertex set of one of the trees of $F$. In this situation, condition (ii) holds. Assume, next, that neither condition (i) nor condition (ii) hold. Then the parent $x_{q}$ of $f$ in $T$ has at most one child in $F$,

```
Algorithm 4 find_f \((F)\)
input: a forest \(F\) (a representative tree with two missing edges)
output: the fixed element \(f \leq 2 n+1\)
1. if \(F\) contains a large vertex \(x\) having a sibling \(z\), then
    return \(f:=\max \{x, z\}\)
2. for each large vertex \(x\) of \(F\) satisfying \(N_{F}(x) \neq \emptyset\)
    for each small \(y \in N_{F}(x)\)
        \(Y^{\prime} \leftarrow\{x-n, x-n+1, \ldots, n\}\)
        if \(\left(N_{F}^{*}(x)=Y^{\prime}\right.\) or \(\left.N_{F}^{*}(x) \subset Y^{\prime}\right)\) and
            \(\left(Y^{\prime} \backslash N_{F}^{+}(x)\right.\) is the vertex set of a tree of \(\left.F\right)\) then
            return \(f:=x\)
```

3. Find the preorder traversals of the 3 trees of $F$, and let $f$ be the unique vertex that is both large and the rightmost element of the preorder traversal of some tree of $F$.
return $f$
whereas $f$ has no children. The latter implies that $f$ is the rightmost vertex of the tree of $F$ containing it. Since $x_{q} \neq f$, we know that no more than two edges have been deleted from $T$, hence no large vertex other than $f$ can be a leaf of $F$. Consequently, condition (iii) holds, completing the proof.

The above theorems lead to an algorithm that efficiently finds the fixed vertex of watermark $G$ (see Algorithm(4). The input is the forest $F$, obtained from the representative tree $T$ of $G$ by the removal of two edges. First, the algorithm checks whether $f=2 n+1$. By Theorem 25, it suffices to verify whether $2 n+1$ is a leaf of $F$ and all small vertices are children of $2 n$, except possibly two, which must be isolated vertices. If this is not the case, then the algorithm proceeds to determining $f$ knowing that $f<2 n+1$. Basically, such task consists in checking conditions (i), (ii) and (iii) of Theorem 26, which can be done in a straightforward manner.

Steps 1 and 3 can be computed easily in linear time. As for Step 2, observe that there are at most two large vertices $x$ of $F$ that may satisfy the condition of having only small children. Consequently, the tests in Step 2
apply to at most two candidates $x$, hence the entire algorithm runs in $O(n)$ time.

### 5.3. Determining the root's children

After having identified the fixed vertex of the watermark, we are almost in a position to determine the tree edges that have been removed.

Observe that, when $f=2 n+1$, the task is trivial, since, in this case, by Theorem [22, there can be only one canonical reducible permutation graph $G$ relative to $n$. Such graph is precisely the one with a Type- 1 representative tree $T$, which is unique for each $n>2$ (cf. Property 3 of canonical reducible permutation graphs, in Section (3). By definition, the root-free preorder traversal of a Type- 1 representative tree, when $f=2 n+1$, is $n+1, n+2, \ldots, 2 n, 1,2, \ldots, n, 2 n+1$.

We therefore want to determine the children of $2 n+2$ restricted to the case where $f<2 n+1$. Let $G$ be a watermark, $T$ its representative tree and $F$ the forest obtained from $T$ by the removal of two edges. As usual, $f$ stands for the fixed vertex of $T, X$ is the set of large vertices other than $2 n+2$, and $X_{c}=X \backslash\{f\}$. Finally, denote by $A \subseteq X_{c}$ the subsets of ascending large cyclic vertices of $T$, which we shall refer to simply as the ascending vertices, and denote by $D$ the set $D=X_{c} \backslash A$ of descending large cyclic vertices of $T$, or simply the descending vertices. Given the forest $F$ and its fixed vertex $f$, Algorithm 5 computes the set $A$, which, as we recall from the proof of Theorem [22, corresponds precisely to the children of the root $2 n+2$.

It is easy to conclude that the above algorithm can be implemented in $O(n)$ time. Now we prove its correctness.

Theorem 27. Algorithm 5 correctly computes the set of ascending vertices $A$ of $T$.

Proof: We follow the different conditions that are checked by the algorithm. Assume $F\left[X_{c}\right] \cup\{2 n+2\}$ is connected. Then $N_{T}(2 n+2)=N_{F}(2 n+2)$, implying $A=N_{F}(2 n+2)$. The algorithm is therefore correct if it terminates at Step 1.

Assume $F\left[X_{c}\right] \cup\{2 n+2\}$ is disconnected, but has no isolated vertices. Then either $N_{F}(2 n+2)=N_{F}(2 n+2)$ or the edge $(2 n+2,2 n+1)$ was one of those that might have been removed from $T$. In any of these situations, we can write $A=N_{F}(2 n+2) \cup 2 n+1$, implying that the algorithm is also correct if it terminates at Step 2.

```
Algorithm 5 find_ascending_large_vertices \((F)\)
input: a forest \(F\) (a representative tree with two missing edges)
output: the children \(A\) of the root \(2 n+2\) of the representative tree
1. if \(F\left[X_{c}\right] \cup 2 n+2\) is connected then
return \(A:=N_{F}(2 n+2)\)
2. if \(F\left[X_{c}\right] \cup 2 n+2\) contains no isolated vertices then
return \(A:=N_{F}(2 n+2) \cup 2 n+1\)
```

3. if $F\left[X_{c}\right] \cup 2 n+2$ contains two isolated vertices $x, x^{\prime}$ then return $A:=N_{F}(2 n+2) \cup\left\{x, x^{\prime}\right\}$
4. if $F\left[X_{c}\right] \cup 2 n+2$ contains a unique isolated vertex $x$ then
if $\left|N_{F}^{*}(f)\right|=2 n-f+1$ then
let $y_{r}$ be the rightmost vertex of $N_{F}^{*}(f)$ if $\left|N_{F}(2 n+2)\right|<y_{r}$ then
return $A:=N_{F}(2 n+2) \cup\{x, 2 n+1\}$
else
return $A:=N_{F}(2 n+2)$
else
return $A:=N_{F}(2 n+2) \cup\{x\}$

Assume $F\left[X_{c}\right] \cup\{2 n+2\}$ contains two distinct isolated vertices $x, x^{\prime}$. The only possibility is $x, x^{\prime} \in N_{T}(2 n+2)$. So, the action of constructing $A$ as the union of $x, x^{\prime}$ and $N_{F}(2 n+2)$ assures correctness, whenever the algorithm terminates at Step 3.

The last situation is $F\left[X_{c}\right] \cup\{2 n+2\}$ containing a unique isolated vertex $x$. We consider the following alternatives. If $\left|N_{F}^{+}(f)\right|=2 n-f+1$, it implies that $N_{T}^{*}(f)=N_{F}^{*}(f)$, because the set of descendants of $f$ in $T$ comprises exactly $y_{f_{0}}, y_{f_{0}+1}, \ldots, y_{n}$, The number of such descendants of $f$ is therefore $n-f_{0}+1$, which, by Property 3 of canonical reducible permutation graphs, is equal to $2 n-f+1$. Now, by Theorem [22, $\left|N_{T}(2 n+2)\right|=y_{r}$, where $y_{r}$ is the rightmost vertex of $N_{F}^{*}(f)$. In this situation, $\left|N_{F}(2 n+2)\right|<y_{r}$ implies that $x$ necessarily belongs to $N_{F}(2 n+2)$. In addition, edge $(2 n+2,2 n+1)$ might also have been deleted from $T$, since a single edge deletion suffices to turn $x$ into an isolated vertex. Observe, on the other hand, that isolating a large
vertex which is not a child of $2 n+2$ requires the removal of at least two edges, provided $n>2$. Thus, $A=N_{F}(2 n+2) \cup\{x, 2 n+1\}$, and the algorithm is correct. In case $\left|N_{F}(2 n+2)\right|=y_{r}$, we know that $N_{T}(2 n+2)=N_{F}(2 n+2)$, hence $A=N_{F}(2 n+2)$, ensuring the correctness of the algorithm. Finally, when $\left|N_{F}^{*}(f)\right| \neq 2 n-f+1$, it means some edge inside the subtree rooted at $f$ has been deleted from $T$. In this case, the isolated vertex $x$ is necessarily a child of $2 n+2$ in $T$, implying $A=N_{F}(2 n+2) \cup\{x\}$, and the algorithm is correct.

### 5.4. Retrieving the missing edges

Once we know the set of ascending vertices, it is simple to restore the entire tree $T$. Basically, given sets $A$ and $X_{c}$, we obtain the set of $D$ of descending vertices. Then, by sorting $A$ and $D$ accordingly, we can locate all the large cyclic vertices in $T$, using the model given by Theorem 22, We then place $f$ in $T$, such that its parent $x_{q}$ is smallest cyclic vertex that is larger than $f$. Finally, we place the small vertices. Vertices $\{1,2, \ldots, f-n-1\}$ are all children of $x_{n}$. The remaining small vertices $\{f-n, f-n+1, \ldots, n\}$ are descendants of $f$ and their exact position in $T$ can be obtained as follows. For each $y \in\{f-n, f-n+1, \ldots, n\}$, we find its position in the preorder traversal $P$ of $T$ by determining the large vertex $x$ whose position in the bitonic sequence of the cyclic large vertices is exactly $y$. Then $y$ must be the $x$ th vertex in the root-free preorder traversal of $T$. Finally, the position of $f$ in $P$ is clearly equal to $f$.

The details are given in Algorithm 6, which computes the preorder traversal $P$ of $T \backslash 2 n+2$.

Again, it is straightforward to conclude that Algorithm 6 correctly computes the preorder traversal of $T$ in time $O(n)$. Such procedure assures the complete retrieval of $T$ and therefore we are able to restore the watermark $G$ in full.

### 5.5. A new decoding algorithm

We can now formulate our new decoding algorithm. If the input watermark presents $k \leq 2$ missing edges, the algorithm is able to fix it prior to running the decoding step. The decoding step itself is absolutely straightforward, and relies on the following theorem.

```
Algorithm 6 retrieve_preorder_traversal \(\left(T, f, A, X_{c}\right)\)
input: the fixed vertex \(f\), the set \(A\) of ascending vertices,
    and the set \(X_{c}\) of large cyclic vertices
    in a representative tree \(T\)
output: the preorder traversal \(P\) of \(T\)
```

1. Let $D \leftarrow X_{c} \backslash A$.
2. The initial vertices of $P$ are those of $A$ in ascending order, followed by those of $D$, in descending order.
Now, subsequently place in $P$ the small vertices $1, \ldots, f-n-1$, in this exact order, immediately after the last descending vertex $x_{n} \in D$. Then place $f$ as to immediately follow $f-n-1$.
3. For each small vertex $y \in\{f-n, f-n+1, \ldots, n\}$, let $P[y]$ be the (large) vertex $x$ whose index in $P$ is $y$, and place $y$ at position $x$ in $P$, i.e., satisfying $P[x]=y$.

## 4. return $P$

Theorem 28. Let $\omega$ be a given identifier and $G$ the watermark corresponding to $\omega$. Let $A=x_{1}, \ldots, x_{\ell-1}$ be the ascending sequence of children of $2 n+2$, in the representative tree $T$ of $G$, that are different from $2 n+1$. Then

$$
\omega=\sum_{i=1}^{\ell-1} 2^{2 n-x_{i}}
$$

Proof: The children of $2 n+2$ in $T$ are the vertices $x_{i}$ of $G$ which are the tail of some tree edge of $G$ pointing to $2 n+2$. From Property 8 of canonical reducible permutation graphs, such vertices $x_{i} \neq 2 n+1$ are precisely those satisfying $x_{i}=n+z_{i}$, where $z_{i}$ is the index of a digit 1 in the binary representation $B$ of $\omega$. The summation yielding $\omega$ can now be easily checked, since the relative value of a digit 1 placed at position $z_{i}$ is $2^{n-z_{i}}=2^{n-\left(x_{i}-n\right)}=2^{2 n-x_{i}}$.

As a consequence of the above theorem, whenever the input watermark has not been tampered with, the proposed Algorithm 7 is able to retrieve the encoded identifier in a very simple way. Note that, in this case, it is not

```
Algorithm 7 decode \((G)\)
input: a watermark \(G\) with \(2 n+3\) vertices and \(0 \leq k \leq 2\) missing edges
output: the identifier \(\omega\) encoded by \(G\)
```

1. Let $k \leftarrow|E(G)|-(4 n+3)$.
2. If $k>2$, report the occurrence of $k$ edge removals and halt.
3. If $0<k \leq 2$, proceed to the reconstitution of the watermark (see Section 5, Algorithms 116).
4. Calculate and return the identifier $\omega$ as indicated by Theorem 28 .
even necessary to obtain the representative tree of the watermark, since the set $A$ can be determined as $A=N_{G}^{-}(2 n+2)$.

Theorem 29. Algorithm $\backslash$ retrieves the correct identifier, encoded in a watermark with up to two missing edges, in linear time.

Proof: Since the final step of the algorithm clearly runs in linear time, its overall time complexity relies on the fact that Algorithms $1 \times 6$ run in linear time themselves, as proved earlier in the text. The correctness of the algorithm follows from the fact that those procedures guarantee the reconstruction of the original watermark when $k \leq 2$ edges have been removed, and from the correctness of Theorem 28,

Corollary 30. Distortive attacks in the form of $k$ edge modifications (insertions/deletions) against canonical reducible permutation graphs $G$, with $|V(G)|=2 n+3, n>2$, can be detected in polynomial time, if $k \leq 5$, and also recovered from, if $k \leq 2$. Such bounds are tight.

Proof: From Theorem [29, we know that, for $n>2$, there are no two watermarks $G_{1}, G_{2}$, with $\left|V\left(G_{1}\right)\right|=\left|V\left(G_{2}\right)\right|=2 n+3$, such that $\mid E\left(G_{1}\right) \backslash$ $E\left(G_{2}\right) \mid \leq 2$, otherwise it would not always be possible to recover from the removal of up to two edges. Thus, for $n>2$, any two canonical permutation graphs $G_{1}, G_{2}$ satisfy

$$
\begin{equation*}
\left|E\left(G_{1}\right) \backslash E\left(G_{2}\right)\right|=\left|E\left(G_{2}\right) \backslash E\left(G_{1}\right)\right| \geq 3 \tag{1}
\end{equation*}
$$

hence $G_{1}$ cannot be transformed into $G_{2}$ by less than 6 edge modifications. Since the class of canonical permutation graphs can be recognized in polynomial-time in light of the characterization given in Theorem 22, and since any number $k \leq 5$ of edge modifications made to a graph $G$ of the class produces a graph $G^{\prime}$ that does not belong to the class, all distortive attacks of such magnitude $(k \leq 5)$ can be detected. Now, for $k=2$, we have three possibilities:
(i) two edges were removed;
(ii) two edges were inserted;
(iii) one edge was removed and one edge was inserted.

If case (i) applies, Theorem 29 guarantees that the original graph can be successfully restored. If case (ii) or case (iii) apply, then a simple algorithm in which all possible sets of two edge modifications are attempted against the damaged graph $G^{\prime}$ suffices to prove that the original graph $G$ can be restored in polynomial time, since, as we already know, exactly one such set shall turn $G^{\prime}$ into a canonical reducible permutation graph. The case $k=1$ is simpler and can be tackled in analogous manner.

It remains to show that such bounds are tight. We present a pair of canonical permutation graphs $G_{1}, G_{2}$, with $\left|V\left(G_{1}\right)\right|=\left|V\left(G_{2}\right)\right|=2 n+3, n>2$, such that inequation (1) holds with equality. We remark that there are many such pairs, and the following is but an example. Let $G_{1}, G_{2}$ be the watermarks relative to identifiers $\omega_{1}=8, \omega_{2}=9$, respectively. Their edge sets are such that $E\left(G_{1}\right) \backslash\{(2,3),(7,8),(8,9)\}=E\left(G_{2}\right) \backslash\{(2,4),(7,9),(8,10)\}$, completing the proof.

## 6. Polynomial-time decoding ( $k$ missing edges)

The linear-time recognition of the class of canonical reducible permutation graphs, wrapped up in the form of Corollary 4, allows the construction of a polynomial-time algorithm to recover watermarks which have been deprived of $k$ edges, for arbitrary values of $k$. The proposed algorithm is formally robust [18], since it manages to repair a damaged watermark $G^{\prime}$ whenever such a thing is possible; otherwise, rather than producing a misled result, it shows that $G^{\prime}$ does not belong to the family of damaged watermarks that can possibly be recovered. As a certificate for this latter case, it outputs two or more watermarks that may become isomorphic to $G^{\prime}$ through the removal
of exactly $k$ of their edges, thus proving that the intended restore is not at all possible.

Let $G$ be a watermark and $G^{\prime}$ the graph obtained from $G$ when a certain subset of $k$ edges are removed. The idea is simple. The algorithm attempts the addition to $E\left(G^{\prime}\right)$ of each and every $k$-subset of non-edges of $G^{\prime}$, one subset at a time. After each attempt, it checks whether a valid watermark (i.e., a canonical reducible permutation graph) was produced. If, after trying all subsets, only one graph was recognized as such, then the decoding was successful. Otherwise, it displays a set containing all watermark candidates.

Since $\left|V\left(G^{\prime}\right)\right|=2 n+3$ and $\left|E\left(G^{\prime}\right)\right|=4 n+3-k$, the number of $k$-subsets of non-edges of $G^{\prime}$ is

$$
\binom{\binom{2 n+3}{2}-(4 n+3-k)}{k}=\mathcal{O}\left(n^{2 k}\right)
$$

Thus, considering the effort of running the recognition algorithm for each one of these watermark candidates, the algorithm runs in overall $\mathcal{O}\left(n^{2 k}\right) \cdot \mathcal{O}(n)=$ $\mathcal{O}\left(n^{2 k+1}\right)$ time.

The aforementioned formulation considers that all non-edges of $G^{\prime}$ could be an edge of the original watermark. However, owing to the particular structure of canonical reducible permutation graphs, relatively few among those non-edges do really stand a chance of belonging to $G$. More precisely, every vertex $v$ of $G$ has out-degree at most 2 , hence $v$ must be the tail endpoint of a most $2-\left|N_{G^{\prime}}^{+}(v)\right|$ edges. The multiset $M^{*}$ of all candidates to being the tail of a missing edge has therefore

$$
\begin{aligned}
|M *| & =\sum_{v \in V\left(G^{\prime}\right)}\left(2-\left|N_{G^{\prime}}^{+}(v)\right|\right) \\
& =2 \cdot\left|V\left(G^{\prime}\right)\right|-\sum_{v \in V\left(G^{\prime}\right)}\left|N_{G^{\prime}}^{+}(v)\right| \\
& =2 \cdot\left|V\left(G^{\prime}\right)\right|-\left|E\left(G^{\prime}\right)\right| \\
& =2 \cdot(2 n+3)-(4 n+3-k) \\
& =k+3
\end{aligned}
$$

elements (not necessarily distinct), and therefore the $k$ missing edge tails may be chosen in $\binom{\left|M^{*}\right|}{k}=\mathcal{O}\left(k^{3}\right)$ different ways. For each $k$-subset of $M^{*}$, the algorithm must choose the head corresponding to each tail, which can be done in $\mathcal{O}(n)$ ways per edge, for an overall $\mathcal{O}\left(k^{3} n^{k}\right)$ number of $k$-subsets
of non-edges that shall be tentatively added to $G^{\prime}$. With the $\mathcal{O}(n)$ running time of the recognition algorithm for each such attempt, we complexity of the whole decoding algorithm is an overall $\mathcal{O}\left(k^{3} n^{k+1}\right)$.

As a matter of fact, it is still possible to cut a whole $\mathcal{O}\left(k^{3}\right)$ factor from that asymptotical complexity, if the labels of the vertices are known. To assume that the labels are known is reasonable in many situations, since each vertex corresponds to a block in the CFG of the software, and, by construction, the watermark graph possesses a Hamiltonian path $2 n+2,2 n+1, \ldots, 0$ that corresponds, in the CFG, to a chunk of subsequent blocks. If that is the case, then we know the out-degree, in $G$, of all watermark vertices (the tail and the head of the Hamiltonian path have out-degrees 1 and 0 , respectively; all other vertices have out-degree 2) and, consequently, the tails of all missing edges. By running the linear-time recognition algorithm on each possible choice of heads, the decoding algorithm has an overall $\mathcal{O}\left(n^{k+1}\right)$ time complexity.

## 7. Proofs of the properties from Section 3

We had postponed the proofs of the properties stated in Section 3 to avoid an overhead of technical pages too early in the paper. We now present the full proofs.

Proof of Property 1. When read from right to left, the $n$ rightmost elements in $P_{b}$ correspond to the $n$ first elements in $Z_{1}$, i.e the $n$ first indexes, in $B^{*}$, where a digit 1 is located. Since $B^{*}$ starts with a sequence of $n$ contiguous 1's, the property ensues.

Proof of Property 2. In $B^{*}$, digits with indexes $1,2, \ldots, n$ are all 1 , by construction. Since the $n$ rightmost elements in $P_{b}$ (i.e., elements indexed $n+2 \leq i \leq n^{*}$ in $P_{b}$ ) correspond to the first $n$ elements in $Y$, and therefore to the first $n$ indexes of 1 's in $B^{*}$, those will always be precisely the elements of set $S=\{1,2, \ldots, n\}$. In other words, if $s \in S$, then $s$ will have index $n^{*}-s+1>n+1$ in $P_{b}$. By the time the elements of $P_{b}$ are gathered together in pairs with views to defining their placement in $P_{s}$, element $s$ will be paired with element $q$ whose index is $n^{*}-\left(n^{*}-s+1\right)+1=s$. Because $s \leq n$, such $q$ clearly does not belong to $S$, hence $q>n$. Now, because $s$ will be assigned index $q$ in $P_{s}$, the element with index $s$ in $P_{s}$ will be its pair $q>n$, concluding the proof.

Proof of Property 3. The bitonic permutation $P_{b}$ is assembled in such a way that its $(n+1)$ th element $f=b_{n+1}$ is either:
(i) the $(n+1)$ th element of $Z_{0}$, in case $B^{*}$ has at least $n+1$ digits 0 ; or
(ii) the $(n+1)$ th element of $Z_{1}$, otherwise.

By construction, the number of 0 's in $B^{*}$ is one unit greater than the number of 1's in $B$.

If (i) holds, then $B$ corresponds to an identifier $w$ that is the predecessor of a power of 2 , implying all $n$ digits of $B$ are 1 's. If that is the case, then the desired property follows immediately, once the $(n+1)$ th element of $Z_{0}$ will be the index of the $(n+1)$ th - i.e., the last - digit 0 in $B^{*}$. Such index is, by construction, $n^{*}$.

If (ii) holds, then $f$ is the index of the $(n+1)$ th digit 1 in $B^{*}$. By construction, the $n$ first digits 1 in $B^{*}$ occupy positions with indexes $1, \ldots, n$, and the $(n+1)$ th digit 1 in $B^{*}$ corresponds to the first digit 1 in the one's complement of $B$. Since that digit has index $f_{0}$ in the one's complement of $B$, and there are in $B^{*}$ exactly $n$ digits to the left of the one's complement of $B$, the property follows.

Proof of Property 居 From the construction of $P_{s}$ and Property 1, it follows that the elements that occupy positions with indexes $1,2, \ldots, n$ in $P_{s}$ are the first $n$ elements in $P_{b}$. It just occurs that the first $n_{1}+1$ numbers in $P_{b}$ are the elements of $Z_{0}$, i.e., the indexes of 0 's in $B^{*}$. Now, the last digit in $B^{*}$ - the one indexed $n^{*}$ - is always a 0 . Besides that 0 , the other digits 0 in $B^{*}$ have indexes $z=n+d$, where each $d$ is the index of a digit 1 in $B$ (the original binary representation of the identifier $\omega$ ). While the first digit in $B$ is always 1 , it is also true that:
(i) the $f_{0}-1$ first digits in $B$ constitute a seamless sequence of 1 's, in case there is at least one 0 in $B$; or
(ii) all $n$ digits of $B$ are 1's, in which case $\omega$ is the predecessor of a power of 2 .

Whichever the case, Property 3 allows us to state that there is a sequence of $f-n-1$ digits 1 in $B$ starting at the first digit of $B$. Such sequence will show up, in $B^{*}$, starting at index $n+1$, in such a way that the $f-n-1$
first elements of $Z_{0}$ will be $n+1, n+2, \ldots, n+(f-n-1)=f-1$. Those elements, as we have seen, will be precisely the first numbers in $P_{b}$. Because there are no more than $n$ such elements, they will be paired against elements $1,2, \ldots, f-n-1 \leq n$ (located from the right end of $P_{b}$ leftwards) in order to determine their placement in $P_{s}$, and the property follows.

Proof of Property 5. If the identifier $\omega$ is not the predecessor of a power of 2 , then its binary representation $B$, whose first digit is always a 1 , contains some digit 0 . In light of this, Property 3 implies $f \geq n+2$ for all integers $\omega$, and the first equality now follows from Property 4. The second equality is granted by the self-invertibility of $P_{s}$, whereby $s_{j}=u \Longleftrightarrow s_{u}=j$.

Proof of Property 6. First, note that $f \neq n^{*}$ corresponds to the case where the identifier $\omega$ is not the predecessor of a power of 2 , i.e., $n_{1}<n$. Because the sequence $Z_{0}$ has exactly $n_{1}+1$ elements, the last of which being the index $n^{*}$ of the rightmost digit in $B^{*}$, element $n^{*}$ will always be assigned index $n_{1}+1$ in $P_{b}$. As we have seen in the proof of Property 4, for $i \leq n$, the $i$ th element in $P_{b}$ will also be the $i$ th element in $P_{s}$, for it will be paired against element $i$, indexed $n^{*}-i+1$ in $P_{b}$ (due to the starting sequence of $n$ digits 1 in $B^{*}$ ). That being said, element $n^{*}$, indexed $n_{1}+1 \leq n$ in $P_{b}$, will have index $n_{1}+1$ in $P_{s}$ as well. If $f=n^{*}$, then the definition of $f$ verifies the property trivially.

Proof of Property 7. We employ again the fact, noted for the first time in the proof of Property 4, that the subsequence consisting of the first $n$ elements in $P_{s}$ and the subsequence consisting of the first $n$ elements in $P_{b}$ are one and the same. Since $P_{b}$ is bitonic, whatever subsequence of $P_{b}$ is bitonic too, particularly the one containing its first $n$ elements. By Property 5 , the central element $s_{n+1}$ of $P_{s}$ is always equal to 1 , therefore the bitonic property of the subsequence consisting of the leftmost elements of $P_{s}$ will not be broken after its length has grown from $n$ to $n+1$, that is, after element $s_{n+1}=1$ has been appended to it.

Proof of Property [8. The first $n_{1}+1$ elements of the bitonic permutation $P_{b}$ are the elements of $Z_{0}$, corresponding to the indexes of 0 's in the extended binary $B^{*}$ (which consists, we recall, of $n$ digits 1 , followed by the one's complement of the binary representation $B$ of the identifier $\omega$ encoded by $G$, followed by a single digit 0 ). Those elements constitute the ascending prefix $A=n+z_{1}, n+z_{2}, \ldots, n+z_{n_{1}}, 2 n+1$, where, for $i \in\left\{1, \ldots, n_{1}\right\}, z_{i}$ is the index of a digit 1 in $B$. From the proof of Property 4. we know that, for $i \leq n$, the $i$ th element in $P_{b}$ will also be the $i$ th element in the self-inverting permutation $P_{s}$. Since $n_{1} \leq n$, we have that the $n_{1}$ first elements of $P_{s}$ are precisely the $n_{1}$ first elements of $A$, hence the tree edge tailed at each of those elements must point, by construction, to $2 n+2$. It remains to show that no element $u \notin A \cup\{2 n+1\}$ is the tail of a tree edge pointing to $2 n+2$. But this comes easily from the fact that, by Property 6, the $\left(n_{1}+1\right)$ th element in $P_{s}$ is $2 n+1$. Since all vertices $u$ with indexes $i>n_{1}+1$ in $P_{s}$ are certainly smaller than $2 n+1$, they can only be the tail of tree edges pointing to vertices $q(u) \leq 2 n+1$, and the proof is complete.

Proof of Property 9. Both items are trivially verified, since, by construction, every tree edge $(u, k) \in G$ is such that either $k>u$ is the element that is closest to $u$ and to the left of $u$ in $P_{s}$, or $k=2 n+2$.

## 8. Proof of Theorem [24]

Let $\mathcal{G}_{k}$ be the set of all canonical reducible permutation graphs with $k$ edges missing. When an element $G^{\prime}$ of $\mathcal{G}_{k}$ is the input of plug_next_subpath $\left(G^{\prime}, \emptyset, \mathcal{H}\right)$, its output is clearly a Hamiltonian path of some graph $G$ such that $V(G)=V\left(G^{\prime}\right)$ and $E(G) E\left(G^{\prime}\right) \leq k$. Thus, when a canonical reducible permutation $G$ minus two edges is passed to Algorithm [1, the path $H$ it returns is the Hamiltonian path of some element of $\mathcal{G}_{2}$. We claim such graph can be no other but $G$.

Let $\hat{H}=2 n+2,2 n+1, \ldots, 0$ be the unique Hamiltonian path of $G$. We divide the proof in three cases:
(i) the removed edges were both tree edges of $G$;
(ii) the removed edges were both path edges of $G$;
(iii) the removed edges were one tree edge and one path edge of $G$.
(a)

(b)

(c)

(d)

(e)

(f)


Figure 5: Possible scenarios for the Hamiltonian path $H$ of a damaged watermark $G^{\prime}$. Dashed arrows indicate missing edges. Squares, solid circles and hollow circles represent vertices whose out-degrees in $G^{\prime}$ are, respectively, 0,1 and 2 . Three hollow circles close together followed by a broken arrow (with a tilde in the middle) indicate subpaths of zero or more edges. Each big rectangle encloses a maximal undamaged subpath of $H$, which corresponds to a maximal backward-unbifurcated path to $s \in V(G)$, i.e., an $\{s\}$-bup.

When two tree edges are missing. The easiest case is (i), as illustrated in Figure 5(a). If only tree edges were removed, then the Hamiltonian path of $G$ is undamaged. Starting from the only vertex with out-degree zero in $G^{\prime}$, namely vertex 0 , plug_next_subpath $\left(G^{\prime}, \emptyset, \mathcal{H}\right)$ outputs $\hat{H}$ at once, never making a single recursive call. Since $\hat{H}$ obviously produces the correct labeling of vertices of $G$, it is validated uneventfully by validate_labels $\left(G^{\prime}, \hat{H}\right)$ and returned by the algorithm.

When two path edges are missing. Suppose now that (ii) is the case. Since no tree edges were removed, the only vertex with out-degree zero in $G^{\prime}$ is
vertex 0 , unless the path edge whose tail is $2 n+2$ was one of the removed edges (we recall that $2 n+2$ has degree 1 in $G$ ). We therefore analyze two subcases, according to whether or not $(2 n+2,2 n+1)$ was removed from $E(G)$.

For the first subcase, suppose the removed path edges were $(2 n+2,2 n+1)$ and ( $a, a-1$ ) for some $a \geq 1$, as in Figure 5(b). Although vertex $2 n+2$ has degree zero, its in-degree is greater than 1 in $G^{\prime} \backslash \emptyset=G^{\prime}$, and therefore no $\{2 n+2\}$-bup is produced in the main call to plug_next_subpath. Thus, the only partial path the algorithm produces, starting from vertex 0 in backwards fashion, is $Q^{\prime}=a-1, a-2, \ldots, 0$. Because now $a-1$ has no in-neighbors in $G-V\left(Q^{\prime}\right)$, it recurses to find possible extensions for $Q^{\prime}$. The only vertices with degree 0 or 1 in $G^{\prime} \backslash V\left(Q^{\prime}\right)$ are now $2 n+2$ and $a$. However, any $\{2 n+2\}$-bup $Q^{\prime \prime}$ that may be found will constitute an $H=Q^{\prime \prime} \| Q$ path that will necessarily fail the ensuing validation. This is due to the fact that $H[2 n+2]$ will be a vertex other than $\hat{H}[2 n+2]=2 n+2$, an out-neighbor of $n+1$ by Theorem 22 and the first condition tested by validate_labels $\left(G^{\prime}, H\right)$ cannot be met. Because of this, the partial path $Q^{\prime}$ can only be extended by an $\{a\}$-bup, which can be no other but $2 n+1, \ldots, a$, and the remaining vertex $2 n+2$ will be concatenated during the next recursive call, completing the Hamiltonian path $2 n+2,2 n+1, \ldots, a, a-1, \ldots, 0=\hat{H}$, as desired.

The second subcase is the one in which the path edge $(2 n+2,2 n+1)$ was not removed. Suppose the missing path edges are $(a, a-1)$ and $(b, b-1)$, with $a<b$, as illustrated in Figure 5(b). The first, rightmost subpath located by the algorithm can only be the unique $\{0\}$-bup, namely $Q^{\prime}=a-1, a-2, \ldots, 0$. Now there are three vertices whose degree are less than or equal to one: $2 n+2, b$ and $a$.

When the algorithm considers $\{2 n+2\}$-bups during the recursive call to plug_next_subpath $\left(G^{\prime}-V\left(Q^{\prime}\right), Q^{\prime}| | \emptyset, \mathcal{H}\right)$, whichever ensuing Hamiltonian path candidate $H$ it produces will necessarily be discarded. Indeed, if $n+1 \geq a$, then no bup is even produced because the in-degree of $2 n+2$ in $G^{\prime}-V\left(Q^{\prime}\right)$ is at least 2 by the existence of tree edges $(n+1,2 n+2)$ and $(2 n+1,2 n+2)$; and, if $n+1<a$, then $H[n+1]$ is vertex $\hat{H}[n+1]=n+1$ itself and, because its out-degree in $G^{\prime}$ is already 2 , conditions (1) and (2) checked by validate_labels $\left(G^{\prime}, H\right)$ cannot both be met.

When the algorithm considers a $\{b\}$-bup, whichever ensuing Hamiltonian path candidate $H$ it comes up with will also be discarded. Indeed, because the subpath $2 n+2,2 n+1, \ldots, b$ of $\hat{H}$ is intact, vertex $2 n+2$ will be brought into the $\{b\}$-bup before $v$ does, for all $b-1 \geq v \geq a$, hence $H[2 n+2] \neq 2 n+2$.

Moreover, because in particular the tree edge whose tail is $b-1$, say $(b-1, w)$, was not removed, and $w \geq b$, the only possible value for $w$ is $2 n+2$, otherwise there would be a vertex $z \in\{2 n+1,2 n, \ldots, b\}$ with in-degree greater than 1 in the subgraph of $G^{\prime}$ induced by $z$ and by the vertices to the left of $z$ in $H$, which is a contradiction because such path would have been discarded in the last line of Algorithm [2. Thus, vertex $b-1$ is an in-neighbor of $2 n+2$ which was not added to the path before $2 n+2$ was added. If a backward bifurcation has not arisen, then it is only possible that $b-1$ is precisely the vertex to the left of $2 n+2$ in $H$. Repeating the same argumentbased on the fact that the tree edge whose tail is $v$ has not been removedfor all $b-2 \geq v \geq a$, we can infer that the only possible Hamiltonian path candidate produced by the concatenation of a $\{b\}$-bup to the left of $Q^{\prime}$ is $H=a, a+1, \ldots, b-1,2 n+2,2 n+1, \ldots, b, a-1, a-2, \ldots 0$. Now, condition (1) in validate_labels enforces that $H[n+1]$ is the tail of a tree edge pointing to $H[2 n+2]=a$. However, because $H[n+1]>a$, such edge cannot be an actual tree edge of the original graph $G$, hence it must be a path edge. Since the only path edge with head $a$ in $G$ is $a+1$, it follows that $H[n+1]=a+1$. And here we shall have a contradiction, since $a+1$ is the second vertex, left to right, in $H$ (i.e., $H[2 n+1]=n+1$ ), unless $a=b-1$. However, if $a=b-1$, then $H[n+1]=b$, and the existence of edge $(b, a)=(H[n+1], H[2 n+2])$ is necessary to meet condition (1) in the validation procedure. But $(b, a)=(b, b-1)$ is one of the removed edges, therefore it must be reinserted. Condition (2), on its turn, requires that an outgoing edge is added to $2 n+2$ (whose degree is 1 and whose index $i$ in $H$ satisfies $2 n+1 \geq i \geq 1$ ). Along with the plausible path edge ( $b, a-1$ ), which was required to concatenate the $\{b\}$-bup to the left of $Q^{\prime}$, we have a total of 3 new edges, thus violating condition (3).

Finally, when the algorithm considers $\{a\}$-bups, it necessarily produces the subpath $Q^{\prime \prime}=b-1, \ldots, a$, which is concatenated to $Q^{\prime}$, and, because $\{2 n+2\}$-bups cannot possibly yield a valid prefix to $Q^{\prime \prime}| | Q^{\prime}$, the last recursive can only produce the $\{b\}$-bup $2 n+2, \ldots, b$, which completes the reconstitution of $\hat{H}$.

When a tree edge and a path edge are missing. We focus on the the final case (iii), where one path edge and one tree edge were removed. We now consider three subcases separately. In the first one, both the path edge and the tree edge that were removed share the same tail endpoint. In the second one, the tails of the removed edges are distinct. The third case is actually
a special case of the second one, when the tail of the removed path edge is vertex $2 n+2$.

For the first subcase, illustrated in Figure 5 (d), say both removed edges have tail $a \in V\left(G^{\prime}\right)$. In this case, vertex $a$ presents degree zero, just like vertex 0 itself. Any attempts to build a Hamiltonian path $H$ whose suffix is an $\{a\}$-bup, however, shall not succeed. Since vertex $2 n+2$ will be brought into $H$ before vertex 0 does, and because $a$ will be the rightmost vertex in $H$ (i.e., $H[0]=a$ ), a tree edge leaving $2 n+2$ is necessary to satisfy condition (2) of validate_labels $\left(G^{\prime}, H\right)$. But vertex 0 appears with index $i>0$ in $H$, and therefore a plausible path edge must be inserted with 0 as its tail. If the index of 0 is not $2 n+2$, then a tree edge leaving 0 is also called for. If, on the other hand, the index of 0 is $2 n+2$, then, among the two tree edges reaching $H[2 n+2]=0$ that are required by condition (1) of the validation procedure, at least one of them is still missing. In both cases, condition (3) is violated.

The second case is the one depicted in Figure 5(e), where a path edge ( $a, a-1$ ), with $1<a \leq 2 n+1$, and a tree edge $(b, v)$, with $v>b$, were removed. Procedure plug_next_subpath starts by gathering the maximal backward-unbifurcated path $Q^{\prime}$ whose head is 0 , the only vertex with degree zero in $G^{\prime}$. The leftmost vertex of such $\{0\}$-bup is vertex $a-1$, the first vertex whose in-degree is zero in the subgraph of $G^{\prime}$ induced by vertices not in $Q^{\prime}$, and hence $Q^{\prime}=a-1, a-2, \ldots, 0$. Now three vertices have out-degree less than or equal to one: $2 n+2, b$ and $a$.

When the algorithm picks $2 n+2$ as a possible continuation of the backward path under construction, the index of $2 n+2$ in $H$ will be $a$. By Theorem [22, vertex $2 n+1$ is always a child of the root $2 n+2$ in the representative tree $T$ of a canonical reducible permutation graph $G$, and, by Property 8, the number of children $v \leq 2 n$ of $2 n+2$ in $T$ corresponds to the number $n_{1}$ of digits 1 in the binary representation $B$ of the identifier $\omega$ encoded by $G$. As a consequence, the in-degree of $2 n+2$ in $G$ is $n_{1}+1$. We now tackle two distinct situations. In the first one, $a \leq n+1$, whereas in the second one $a>n+1$. If $a \leq n+1$, then the in-degree of $2 n+2$ in $G^{\prime}-V\left(Q^{\prime}\right)$ is the same as in $G^{\prime}$ (i.e., $n_{1}$ ), since all in-neighbors of $2 n+2$ belong to $\{n+1, \ldots, 2 n+1\}$ by the same Theorem [22, Because, along with the path edge $(a, a-1)$, only one tree edge was removed from $G$ to obtain $G^{\prime}$, the indegree of $2 n+2$ in $G^{\prime}$ is at least $n_{1}+1-1=n_{1}$. As a consequence, a backward bifurcation would be noticed on $2 n+2$ unless $n_{1}=1$ and the tail $b$ of the removed tree edge is one of the in-neighbors of $2 n+2$, which in this case are $n+1$ and $2 n+1$. If $b=n+1$, then the tree edge $e=(2 n+1,2 n+2)$
is intact, and the only possible placement of vertex $2 n+1$ in $H$ is at the position immediately to the left of $2 n+2$, so that $e$ functions as a path edge of $H$. Assuming there was no backward bifurcation on $2 n+2$ (which would have caused the path $H$ to be discarded), the only possible tree edge leaving $2 n$ is $(2 n, 2 n+1)$, hence $2 n$ must be placed to the left of $2 n+1$ in $H$. Assuming, similarly, that no backward bifurcation occurred on $2 n+1$, the only possible tree edge leaving $2 n-1$ is $(2 n-1,2 n)$, and so on. This reasoning must continue until finally $a$ is concatenated at the very first position of $H$, yielding $H=a, a+1, \ldots, 2 n+2, a-1, a-2, \ldots, 0$. Now, condition (1) of the validation procedure requires that $H[n+1]$ and $H[2 n+1]$ are in-neighbors of $H[2 n+2]=a$. However, this requirement and condition (2) cannot both be met without violating condition (3), because, since those two vertices $H[n+1]$ and $H[2 n+1]$ are not in $Q^{\prime}$, they are certainly greater than $a$, but there is only one vertex in $G$ which is greater than $a$ and is an in-neighbor of $a$, namely $a+1$. Therefore an extra tree edge is required, but one extra edge is also required by condition (2) - a tree edge leaving $b$-and the plausible path edge $(2 n+2, a-1)$ had already been inserted, which breaks condition (3). We are left with the possibility that the tail of the removed tree edge was $b=2 n+1$. In this case, the tree edge $(n+1,2 n+2)$ is intact, and the vertex immediately to the left of $2 n+2$ in $H$ must be $n+1$. Now, since path edge $(n+2, n+1)$ is not the missing one by hypothesis, vertex $n+2$ must be immediately to the left of $n+1$ in $H$, and, since path edge $(n+3, n+2)$ is not the missing one, vertex $n+3$ must appear immediately to the left of $n+2$, and so on, until $b=2 n+1$ is concatenated at the first position of $H$, yielding $H=2 n+1,2 n, \ldots, a, 2 n+2, a-1, a-2, \ldots, 0$. To satisfy condition (1) of the validation, vertex $H[n+1]$ must be an in-neighbor of $H[2 n+2]=2 n+1$. But, because $n_{1}=1$ ( $\omega$ is a power of 2 ), the root of its Type- 2 representative tree has only two children, which allows item (iii) in Definition 20 to assure that $2 n+1$ has only one child, and this child is not $n+1$, by item (i) of that same definition. Thus, the tree edge $(H[n+1], 2 n+1)$ must be added to satisfy condition (1) of validate_labels $\left(G^{\prime}, H\right)$, and the only vertices with out-degree 1 in $G^{\prime}$ were $b$, which is $2 n+1$ itself, $2 n+2$, which was already added a plausible path edge connecting it to $a-1$, and $a$. It is therefore only possible that $H[n+1]=a$, that is, the missing path edge is necessarily $(n+1, n)$. And here is where condition (4) of the validation procedure comes into play, enforcing that the root $H[2 n+2]$ presents only two children when $\omega$ is a power of 2 . Since that is not the case for the path $H$ so obtained, as can be easily checked, $H$ is discarded. The second situation is the one in
which $a>n+1$. This one is easy, since now $H[n+1]=n+1$, which is the tail of a tree edge pointing to $2 n+2 \neq H[2 n+2]$, and hence conditions (1) and (2) of the validation cannot both be met, unless such tree edge is precisely the one tree edge that was removed. But that would correspond to the subcase shown in Figure 5(d), which we already tackled.

When the algorithm picks $b$ as the head of the first subpath to extend the $\{0\}$-bup $Q^{\prime}$, all ensuing Hamiltonian path candidates shall be discarded by similar reasons.

Finally, when it considers the sound continuation $a$, all conditions obviously pass and $\hat{H}$ is delivered.

The third-and last-possible situation is the one depicted in Figure5(f), where the removed edges were the path edge $(2 n+2,2 n+1)$ and a tree edge $(b, v)$, with $v>b$. There are two vertices with degree zero: 0 and $2 n+2$. When the call to plug_next_subpath $\left.\left(G^{\prime}, \emptyset, \mathcal{H}\right\}\right)$ picks $2 n+2$ as the rightmost vertex of $Q^{\prime}$, the leftmost vertex of whatever Hamiltonian path $H$ it produces must be either 0 or $b$, the only vertices with out-degree less than 2 in $G^{\prime}$ (part of the second condition verified by validate_labels). Moreover, the root of the representative tree of $G$ must have only two children (which means $n_{1}=1$, or, equivalently, the identifier $\omega$ encoded by $G$ is a power of 2 ), and $b$ must be either $2 n+1$ or $n+1$, so that a backward bifurcation does not take place at the very starting vertex $H[0]=2 n+2$. If $H[2 n+2]=0$, then at least three extra edges are required to put $H$ together and satisfy condition (1) of the validation procedure: a plausible path edge $(H[2 n+2], H[2 n+1])$, and at least two tree edges, namely $(H[2 n+1], H[2 n+2])$ and $(H[n+1], H[2 n+2])$. But then, of course, condition (3) is violated. If $H[2 n+2]=b=n+1$, then the vertex immediately to the left of $H[0]=2 n+2$ in $H$ must be $H[1]=2 n+1$, and the next vertex right-to-left must be $H[2]=2 n$ and so on, assuming no backward bifurcations took place, until at least vertex $H[n]=n+2$. To put it differently, the $\{2 n+2\}$-bup $Q^{\prime}$ considered initially by the algorithm contains (not necessarily properly, depending on whether there was a tree edge pointing to $n+2$ in $G^{\prime}$ ) the suffix $Q^{\prime}=n+2, n+3, \ldots 2 n+2$. Now, no matter which vertex $w$ occupies the $(n+1)$ th position (right-to-left) in $H$, it was certainly not an in-neighbor of $H[2 n+2]=n+1$, because $n+1$ does not have in-neighbors in Type-2 trees (and in Type-1 trees neither, for that matter). If $w \neq 0$, then $w$ has out-degree 2 , and conditions (1) and (2) of the validation procedure cannot both be met. If, on the other hand, $w=0$, then $H$ is the concatenation of $Q^{\prime}$ with the prefix $n+1, n, n-1, \ldots, 0$, an intact subpath of $\hat{H}$. In this case, vertex $H[2 n+1]$ is $n$, a vertex with out-degree
(a)

(d)

(b)

(e)

(c)

(1)
(f)


Figure 6: (a) The watermark $G_{1}$ for identifier $\omega=2 ;(\mathrm{b})$ its representative tree $T_{1}$; (c) the damaged representative tree $T_{1}^{\prime}$ obtained from $T_{1}$ by removing edges $(1,5)$ and $(4,5)$; (d) the watermark $G_{2}$ for identifier $\omega=3$; (e) its representative tree $T_{2}$; (f) the damaged representative tree $T_{2}^{\prime}$ obtained from $T_{2}$ by removing edges $(1,4)$ and $(4,6)$. Note that $T_{1}^{\prime}$ and $T_{2}^{\prime}$ are isomorphic.

2 in $G^{\prime}$ which is not an in-neighbor of $H[2 n+2]=n+1$, and conditions (1) and (2), again, cannot both be met.

The verification of the time complexity is straightforward.

## 9. Final considerations

After characterizing the class of canonical reducible permutation graphs, we formulated a linear-time algorithm which succeeds in retrieving $n$-bit identifiers encoded by such graphs (with $n>2$ ) even if $k \leq 2$ edges are missing. Furthermore, we presented a polynomial-time algorithm to decode Chroni and Nikolopoulos's watermarks [5] with an arbitrary number of missing edges whenever it is possible to do so deterministically.

An implication of the first proposed algorithm is that attacks in the form of $k \leq 5$ general edge modifications (deletions/insertions) can always be detected in polynomial time (for $n>2$ ), since the replacement of no more than two edges in a canonical reducible permutation graph yields a graph that does not belong to the class. A minimum of six edge modifications (three deletions followed by three insertions) is therefore necessary to change any
given watermark into a different, valid watermark. Indeed, the sole example of two canonical reducible permutation graphs which may become isomorphic to one another when each graph is deprived of only two edges occurs when $n=2$, as illustrated in Figure 6. The instances $G_{1}$ and $G_{2}$ correspond to identifiers $\omega_{1}=2$ (binary $B=10$ ) and $\omega_{2}=3$ (binary $B=11$ ), respectively. They become isomorphic to one another when edges $(1,5),(4,5)$ are removed from $G_{1}$ and edges $(1,4),(4,6)$ are removed from $G_{2}$. An interesting open problem is to characterize the maximum sets of identifiers $\Omega(k)$ such that, for all $\omega_{1}, \omega_{2} \in \Omega(k)$, the corresponding watermarks cannot become isomorphic to one another when each one is deprived of $k>2$ edges.

Future research focusing on the development of watermarking schemes resilient to attacks of greater magnitude may consider extending the concept of canonical reducible permutation graphs by allowing permutations with $h$-cycles, with $h>2$, as well as multiple fixed elements.

## References

[1] F.E. Allen, Frances E., Control flow analysis, SIGPLAN Not. 5 (1970), 1-19.
[2] Business Software Alliance, Shadow Market: 2011 BSA Global Software Piracy Study (2012), available at http://globalstudy.bsa.org.
[3] M. Chroni and S.D. Nikolopoulos, Efficient encoding of watermark numbers as reducible permutation graphs, arXiv:1110.1194v1 [cs.DS], 2011.
[4] M. Chroni and S.D. Nikolopoulos, Encoding watermark numbers as cographs using self-inverting permutations, Proc. 12th Int'l Conference on Computer Systems and Technologies, CompSysTech'11, ACM ICPS 578 (2011), 142-148 (Best Paper Award).
[5] M. Chroni and S.D. Nikolopoulos, An efficient graph codec system for software watermarking, Proc. 36th IEEE Conference on Computers, Software and Applications, COMPSAC'12, IEEE Proceedings (2012), 595-600.
[6] M. Chroni and S.D. Nikolopoulos, Multiple encoding of a watermark number into reducible permutation graphs using cotrees, Proc. 13th Int'l Conference on Computer Systems and Technologies, CompSysTech'12, ACM ICPS Proceedings (2012), 118-125.
[7] M. Chroni and S.D. Nikolopoulos, An embedding graph-based model for software watermarking, Proc. 8th Int'l Conference on Intelligent Information Hiding and Multimedia Signal Processing, IIH-MSP' 12, IEEE Proceedings (2012), 261-264.
[8] C. Collberg and C. Thomborson, Software watermarking: models and dynamic embeddings, Proc. 26th ACM SIGPLAN-SIGACT on Principles of Programming Languages, POPL'99 (1999), 311-324.
[9] C. Collberg, S. Kobourov, E. Carter and C. Thomborson, Errorcorrecting graphs for software watermarking, Proc. 29th Workshop on Graph-Theoretic Concepts in Computer Science, WG'03, LNCS 2880 (2003), 156-167.
[10] C. Collberg, C. Thomborson and G. Townsend, Dynamic graph-based software fingerprinting, ACM Transactions on Programming Languages and Systems 29 (2007), 1-67.
[11] C. Collberg, A. Huntwork, E. Carter, G. Townsend and M. Stepp, More on graph theoretic software watermarks: implementation, analysis and attacks, Information and Software Technology 51 (2009), 56-67.
[12] C. Collberg, SandMark: A Tool for the Study of Software Protection Algorithms, http://sandmark.cs.arizona.edu/index.html, last visited 25 Nov 2013.
[13] R. L. Davidson and N. Myhrvold, Method and system for generating and auditing a signature for a computer program, US Patent 5.559.884, Microsoft Corporation (1996).
[14] J. Hamilton and S. Danicic, A Survey of Static Software Watermarking, Proc. World Congress on Internet Security, WorldCIS'11 (2011) 100107.
[15] M.S. Hecht and J.D. Ullman, Flow graph reducibility, SIAM J. Computing 1 (1972), 188-202.
[16] M.S. Hecht and J.D. Ullman, Characterizations of reducible flow graphs, Journal of the ACM 21 (1974), 367-375.
[17] F. Nielson, H.R. Nielson and C. Hankin, Principles of Program Analysis (2004), Springer-Verlag.
[18] V. Raghavan and J. Spinrad, Robust Algorithms for Restricted Domains, Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'01 (2001), 460-467.
[19] I. Schaefer, R. Rabiser, D. Clarke, L. Bettini, D. Benavides, G. Botterweck, A. Pathak, S. Trujillo and K. Villela, Software diversity: state of the art and perspectives, International Journal on Software Tools for Technology Transfer 14 (2012), 477-495.
[20] Y. Su, J. Liu and D. Li, Hiding Signatures in Variable Names, Communications in Computer and Information Science 320 (2013), 333-340.
[21] R.E. Tarjan, Testing flow graph reducibiliy, Journal of Computer and System Sciences 9 (1974), 355-365.
[22] R. Venkatesan, V. Vazirani and S. Sinha, A graph theoretic approach to software watermarking, Proc. 4th International Information Hiding Workshop (2001), 157-168.
[23] R. Venkatesan and V. Vazirani, Technique for producing through watermarking highly tamper-resistant executable code and resulting watermarked code so formed (2006), Microsoft Corporation, US Patent: 7051208.
[24] B.A. Wichmann, A.A. Canning, D.L. Clutterbuck, L.A. Winsborrow, N.J. Ward and D.W.R. Marsh, Industrial perspective on static analysis, Software Engineering Journal 10 (1995), 69-75.
[25] J. Zhu, Y. Liu and K. Yin, A novel dynamic graph software watermark scheme, Proc. 1st Int'l Workshop on Education Technology and Computer Science 3 (2009), 775-780.
[26] W. Zhu, C. Thomborson, F-Y. Wang, A survey of software watermarking, Proc. IEEE Int'l Conference on Intelligence and Security Informatics, ISI'05 (2005), 454-458.


[^0]:    ${ }^{*}$ An extended abstract of this paper was presented at the 39th International Workshop on Graph Theoretic Concepts in Computer Science, WG 2013, and appeared in Lecture Notes in Computer Science 8165 (2013), 50-63.

    Email addresses: lucilabento@ppgi.ufrj.br (Lucila M. S. Bento), drboccardo@inmetro.gov.br (Davidson Boccardo), rcmachado@inmetro.gov.br (Raphael C. S. Machado), vigusmao@dcc.ufrj.br (Vinícius G. Pereira de Sá), jayme@nce.ufrj.br (Jayme Luiz Szwarcfiter)

