# Amortized $\tilde{O}(|V|)$-Delay Algorithm for Listing Chordless Cycles in Undirected Graphs* 

Rui Ferreira ${ }^{1}$, Roberto Grossi ${ }^{2}$, Romeo Rizzi $^{3}$, Gustavo Sacomoto ${ }^{4,5}$, and Marie-France Sagot ${ }^{4,5}$<br>${ }^{1}$ Microsoft Bing, UK<br>${ }^{2}$ Università di Pisa, Italy<br>${ }^{3}$ Università di Verona, Italy<br>${ }^{4}$ INRIA Grenoble Rhône-Alpes, France<br>${ }^{5}$ UMR CNRS 5558 - LBBE, Université Lyon 1, France


#### Abstract

Chordless cycles are very natural structures in undirected graphs, with an important history and distinguished role in graph theory. Motivated also by previous work on the classical problem of listing cycles, we study how to list chordless cycles. The best known solution to list all the $C$ chordless cycles contained in an undirected graph $G=(V, E)$ takes $O\left(|E|^{2}+|E| \cdot C\right)$ time. In this paper we provide an algorithm taking $\tilde{O}(|E|+|V| \cdot C)$ time. We also show how to obtain the same complexity for listing all the $P$ chordless $s t$-paths in $G$ (where $C$ is replaced by $P$ ).


## 1 Introduction

A chordless (induced) cycle $c$ in andirected graph $G$ is a cycle such that the subgraph induced by its vertices contains exactly the edges of $c$. A chordless cycle is called a hole when its length is at least 4 . Similarly, a chordless (induced) path $\pi$ in $G$ is such that the subgraph of $G$ induced by $\pi$ contains exactly the edges of $\pi$. Both chordless cycles and paths are very natural structures in undirected graphs with an important history, appearing in many papers in graph theory related to chordal graphs, perfect graphs and co-graphs (e.g. [11|6|3]), as well as many NP-complete problems involving them (e.g. [2|79]).

In this paper we consider algorithms for listing chordless cycles and st-paths in an undirected graph $G=(V, E)$, with $n=|V|$ vertices and $m=|E|$ edges, motivated by the algorithms for listing cycles and $s t$-paths that have been produced by an active area of research since the early 70s 10131.

In this paper we present an algorithm for listing all the $C$ chordless cycles in an undirected graph $G=(V, E)$ in $\tilde{O}(m+n \cdot C)$ time, hence with an amortized $\tilde{O}(n)$ time delay, where $\tilde{O}(f(n, m))$ is used as a shorthand for $O(f(n, m)$ polylog $n)$. We also show that the same algorithm may be used to list

[^0]all the $P$ chordless st-paths in $\tilde{O}(m+n \cdot P)$ time, hence amortized $\tilde{O}(n)$ time delay.

There are very few algorithms in the literature for listing chordless cycles and/or paths, where some of them have no guaranteed performance 12[16. The most notable and elegant listing algorithm is by Uno [15], with a cost of $O\left(m^{2}+m \cdot C\right)$ time for chordless cycles and $O\left(m^{2}+m \cdot P\right)$ time for chordless st-paths, hence amortized $O(m)$ time delay.

## 2 Preliminaries

Our graphs are finite, undirected, and simple, i.e. without self-loops or parallel edges. Given a graph $G=(V, E)$ with $n=|V|$ vertices and $m=|E|$ edges, our task is to list out fast all its chordless cycles. We hence assume that $G$ is connected. Given $V^{\prime} \subseteq V$, we denote by $E\left\langle V^{\prime}\right\rangle:=\left\{u v \in E \mid u, v \in V^{\prime}\right\}$ the set of those edges which are contained in $V^{\prime}$. A graph $G^{\prime}=\left(V^{\prime}, E^{\prime}\right)$ is called a subgraph of $G$ if $V^{\prime} \subseteq V$ and $E^{\prime} \subseteq E$. The subgraph $G^{\prime}$ is called induced (or chordless) if $E^{\prime}=E\left\langle V^{\prime}\right\rangle$. For any $V^{\prime} \subseteq V$, we denote by $G\left[V^{\prime}\right]:=\left(V^{\prime}, E\left\langle V^{\prime}\right\rangle\right)$ the subgraph of $G$ induced by $V^{\prime}$. Given $e \in E$, we denote by $G \backslash e:=(V, E \backslash\{e\})$ the subgraph obtained from $G$ by deleting the edge $e$. Given $v \in V$, we denote by $G \backslash v:=G[V \backslash\{v\}]$ the subgraph obtained from $G$ by first deleting all the edges incident to $v$, and then removing the isolated vertex $v$. Given a vertex $u \in V$, we denote by $N_{G}(u):=\{v \in V \mid u v \in E\}$ the neighbourhood of $u$, the subscript is omitted whenever the graph is clear from the context.

A cycle is a connected graph in which every vertex has degree 2. A path is a connected graph in which every vertex has degree 2 except for two degree- 1 vertices, $s$ and $t$, called the endvertices of the path. This is also called an st-path and denoted by $\pi_{s t}$. Indeed, when building a path from $s$ to $t$ edge after edge, it will be most natural, and more precise, to think like we are orienting the traversed edges. For this reason, we will also write $(u, v)$ for an edge that, when building a path, has been traversed from $u$ to $v$.

A (chordless) path (or cycle) of $G$ is a (chordless) subgraph of $G$ which is a path (or cycle). We denote by $\mathcal{C}(G)$ the set of all chordless cycles in $G$. We denote by $\mathcal{P}(G)$ (by $\left.\mathcal{P}_{s t}(G)\right)$ the set of all chordless paths (st-paths) in $G$. When $s=t$, we get those cycles visiting $s$. We refer to a path $\pi \in \mathcal{P}(G)$ by its natural sequence of vertices or edges. A hole is a chordless cycle of size at least 4. Thus $\mathcal{C}(G)$ comprises holes and triangles. Since there are at most $m n$ triangles, our algorithm can be used to list the holes of $G$ in $\tilde{O}(n)$ time each, with an overall $\tilde{O}\left(m n^{2}\right)$ additive time cost.

Uno [15] proposed an algorithm that lists each chordless cycle in an undirected graph $G=(V, E)$ in $O(m)$ time while using $O(m)$ space. The first step is the following reduction to the problem of enumerating the chordless st-path in a graph $G$. Based on the fact that for any vertex $s \in V$ the chordless cycles in $G \backslash s$ are also chordless cycles in $G$, the algorithm proceeds by listing all chordless cycles passing through $s$; and repeating the process in $G \backslash s$, until the graph is empty. Then, to list all chordless cycles passing through $s$ in $G^{\prime}=G$, the
algorithm follows the approach of listing the chordless paths $s \rightsquigarrow t$ in $G^{\prime} \backslash(s, t)$, for each $t \in N_{G^{\prime}}(s)$; and to avoid duplications, at the end of each iteration the graph is updated to $G^{\prime}=G^{\prime} \backslash t$.

Given a previously computed chordless st-path $\pi=v_{0} v_{1} \ldots v_{l}$, Uno's algorithm identifies the set of vertices $U \subseteq V$ such that each $u \in U$ is adjacent to some $v_{j} \in \pi$, and the edge $\left(v_{j}, u\right)$ is contained in a chordless st-path $\pi^{\prime} \neq \pi$ extending the prefix $\pi_{j}=v_{0} v_{1} \ldots v_{j}$. The algorithm is kick-started by taking a shortest $s t$-path (as a shortest path has the property of also being a chordless path) and employs a recursive strategy of vertex removal to avoid listing the same chordless path multiple times. This ensures that each chordless path is listed once. Uno's algorithm takes $O(m)$ time to compute $U$ and prepare the recursive calls before it either outputs a new path or stops. The total time is therefore $O\left(m^{2}+m \cdot|\mathcal{C}(G)|\right)$.

## 3 Our Approach and Key Ideas

We outline the main ideas which allow us to reduce the amortized cost for a chordless cycle from $O(m)$ to $\tilde{O}(n)$, giving a total $\tilde{O}(m+n \cdot|\mathcal{C}(G)|)$ time to list all the chordless cycles. Our approach relies on a variant of the cleaning operation introduced in [6] to recognize linear balanced matrices and even holes in graphs 4|5.

### 3.1 Certificates for chordless st-path

A listing algorithm usually takes the form of a recursive procedure exploring the space of all solutions. A key idea employed since the first listing papers 10 is to check for the existence of at least one solution before branching, i.e. before partitioning the solution space in subspaces to be assigned to the children. This avoids unproductive recursive calls, i.e. calls that do not list any solution and whose overhead cost could completely dominate the cost of reporting the solutions (e.g. see [14]). In a previous work [1], we stressed the notion of certificate since, in a more refined recursive scheme, passing a certificate of existence as an extra parameter may facilitate the work of the children which may avoid running the existence check: if they have a single child, they could be done by just passing the certificate received as an input or a small adaptation of it. We also saw that more structural facts around the certificate could be useful. For the case of $s t$-paths [1], the certificate is a DFS tree rooted in $s$ and reaching $t$, which contained an st-path and also helped in other ways. Until now, the certificate was itself a solution or explicitely contained one.

Here we try out something new: what if our certificate guarantees the existence of a solution but is not itself a solution? The following fact suggests that the certificate for the existence of a chordless $s t$-path might be just any $s t$-path.

Fact 1 Given two vertices $s, t$ in $G$, there is a chordless st-path in $G$ iff there is an st-path in $G$.

Thus we allow for certificates which are somewhat less refined than actual solutions, in the same spirit that a binary heap demands a less strict and lazy notion of order. This is a new asset of the notion of certificate and opens up new possibilities.

### 3.2 From chordless cycles to chordless st-paths

Uno [15] shows how to reduce listing chordless cycles in a graph to listing chordless st-paths for all edges ( $s, t$ ) chosen in a specific order (see Section 2), which is necessary to avoid duplications in the output. The initialization step for each edge $(s, t)$ takes $O(m)$ time as it requires to find one chordless st-path. This gives the $m^{2}$ term in the total cost of $O\left(m^{2}+m \cdot|\mathcal{C}(G)|\right)$ for chordless cycles.

We observed in Section 3.1 that any st-path will suffice as a starter, as they are our certificates of choice. This makes a difference for the above reduction, since using dynamic graph connectivity algorithms [8, it is possible to maintain a spanning tree in $O$ (polylog$n)$ time per edge deletion, perform connectivity queries in $O(\operatorname{polylog} n)$, and more importantly obtain an st-path in $\tilde{O}(n)$. It is worth noting that it is not known how to obtain a chordless $s t$-path faster than $O(m)$. Hence, we first build the dynamic connectivity structure as preprocessing step. Then, for each edge ( $s, t$ ), in the same order as Uno's reduction, we list the chordless $s t$-paths. Before calling our path listing algorithm for edge ( $s, t$ ), we test if $s$ and $t$ are connected (Fact (1): if so, we call our path listing algorithm, paying $\tilde{O}(n)$ to find one initial st-path; otherwise, we skip the edge $(s, t)$ and take the next in order. As a result, the total initialization cost is $\tilde{O}(m+k n)$ for all edges instead of $O\left(m^{2}\right)$, where $k$ is the number of edges for which we find one initial st-path. Note that $k \leq|\mathcal{C}(G)|$ as each of them surely gives rise to a chordless $s t$-path whence to a distinct chordless cycle. We obtain in this way an $\tilde{O}(m+n \cdot|\mathcal{C}(\underset{\sim}{G})|)$ time algorithm to list chordless cycles, if we can list st-paths in amortized $\tilde{O}(n)$ time each.

### 3.3 Difficulty of cleaning st-paths

Given any st-path, as stated in Fact 1 we can clean it to obtain a chordless stpath in a greedy fashion: start from $u=s$ and iteratively take a neighbour of $u$ that is closest to $t$ along the path. The process stops when $u=t$. The vertices taken in this way form a chordless path. The problem is that the cost of such a greedy traversal of the path is upper bounded by the sum of the degrees of the vertices along it. Unfortunately, this sum could be $\Theta(m)$ in the worst case.

Even worse, this is still true when the initial path is already chordless, as shown in Fig. (1. Consider the complete bipartite clique $K_{r, r}=\left(V_{1} \cup V_{2}, E_{12}\right)$, where $V_{1}=\left\{x_{1}, x_{2}, \ldots, x_{r}\right\}$. Build a new graph $G=(V, E)$ where the vertex set is $V=V_{1} \cup V_{2} \cup\left\{p_{1}, \ldots, p_{r-1}\right\}$ for some new vertices $p_{1}, \ldots, p_{r-1}$, and the edge set is $E=E_{12} \cup\left\{\left(x_{1}, p_{1}\right),\left(p_{1}, x_{2}\right),\left(x_{2}, p_{2}\right), \ldots,\left(x_{r-1}, p_{r-1}\right),\left(p_{r-1}, x_{r}\right)\right\}$. Now, the path $x_{1}, p_{1}, x_{2}, p_{2}, \ldots, x_{r-1}, p_{r-1}, x_{r}$ is chordless but each edge is incident to at least one vertex in that path, so the sum of the degrees is $m=|E|=\Theta\left(r^{2}\right)=$ $\Theta\left(|V|^{2}\right)=\Theta\left(n^{2}\right)$.


Fig. 1. Sum of degrees on chordless path $x_{1}, p_{1}, x_{2}, p_{2}, \ldots, x_{r-1}, p_{r-1}, x_{r}$ is $\Theta(m)$.

What we would like to do: recursively extend a given chordless path $\pi_{s u}$ into a chordless $s t$-path, while maintaining as a certificate an st-path. The recursive extension can be seen as an implicit cleaning of our st-path certificate. Consider a vertex $u$ along a given st-path (our certificate), where initially $u=s$. Our certificate guarantees that there is at least one chordless st-path going through a neighbour of $u$, say $a$. However, exploring all of $u$ 's neighbours would cost too much so we need to proceed more carefully: consider any neighbour $b \neq a$, the following two situations may occur. (1) $a$ and $b$ are both good, meaning that $(u, a)$ and $(u, b)$ are on two distinct chordless st-paths. In this case, the chordless st-paths traversing $(u, a)$ cannot go through $b$ too, as otherwise it would not be chordless (see Remark 1 below), so $b$ should be removed. (2) $b$ is not on any chordless $s t$-path, so it is either disconnected from $t$ or every st-path going through $b$ passes through $a$. In this case, as it will be clear later, we need neither to explore nor to remove $b$.

In other words, we can treat the neighbours of $u$ as described above, and they will not interfere when cleaning the st-path in the next recursive calls since their are either removed (as in case 1) or implicitly cut out (as in case 2). We make this statement more precise below.

### 3.4 Reduced degree property

We introduce a notion of reduced degree with a stronger property in mind. Consider a chordless st-path $\pi_{s t}=v_{0} v_{1} \ldots v_{\ell}$ in the graph $G$, for some integer $\ell>1$, where $v_{0}=s$ and $v_{\ell}=t$. For a vertex $v_{i}$, a neighbour $v \in N\left(v_{i}\right)$ is good if there exists a chordless st-path in $G$ with prefix $v_{0} v_{1} \ldots v_{i} v$ (i.e. it extends $v_{0} v_{1} \ldots v_{i}$ by adding the edge $\left(v_{i}, v\right)$ as illustrated in Fig. (2). We denote by $N^{\text {good }}\left(v_{i}\right) \subseteq N\left(v_{i}\right)$ the set of good neighbours of $v_{i}$, noting that $v_{i+1} \in N^{\text {good }}\left(v_{i}\right)$. For each $v_{i}$, its reduced degree $d_{i}$ is given by the number of non-good neighbours, namely, $d_{i}=\left|\left(N\left(v_{i}\right) \backslash \bigcup_{j \leq i} N^{\text {good }}\left(v_{j}\right)\right) \cup\left\{v_{i+1}\right\}\right|$.

The rationale is that exploring the good neighbours of $v_{i}$ will list further chordless paths while examining its neighbours that are not good is a waste of computation. The reduced degree of $v_{i}$ is actually an upper bound on the number of not-good vertices examined when exploring $v_{i}$ to produce the chordless st-path $\pi_{s t}$ and gives an upper bound on the waste. Lemma 1 below shows that while examining the neighbours of the vertices along a chordless path still takes $O(m)$


Fig. 2. Good neighbours (in red) of vertex $v_{i}$ in $G_{i}$.
time, only $O(n)$ neighbours are a waste while the remaining ones lead to further chordless paths (which is a good argument for amortization).

Lemma 1. For a chordless path $\pi_{s t}$, we have $\sum_{v_{i} \in \pi_{s t}} d_{i} \leq 2 n$, where $d_{i}$ is the reduced degree of $v_{i} \in \pi_{s t}$.

Proof. We will show that each vertex $x$ of $G$ is a non-good neighbour of at most two vertices in $\pi_{s t}$. To this purpose, we prove that if $x$ is a non-good neighbour of both $v_{i}$ and $v_{j}$ then $|i-j| \leq 1$. We choose such three vertices $v_{i}, v_{j}$ and $x$ where the difference $j-i$ is the largest possible and assume by contradiction that $i<j-1$. Thus $v_{i}$ and $v_{j}$ are not adjacent in $\pi_{s t}$, whence $\left(v_{i}, v_{j}\right)$ is not an edge of $G$ since $\pi_{s t}$ is chordless. Also, being non-good, $x \notin \pi_{s t}$. Consider the st-path $\pi^{*}=v_{0} \ldots v_{i} x v_{j} \ldots v_{l}$. Clearly, $\pi^{*}$ contains no repeated vertices and we will prove that $\pi^{*}$ is a chordless $s t$-path, contradicting the fact that $x$ is not a good neighbour of $v_{i}$. The fact that $\pi^{*}$ is chordless follows from the fact that there is no $v_{k} \in \pi^{*}, k \neq i$ and $k \neq j$, such that $\left(v_{k}, x\right)$ is an edge of $G$, otherwise $j-k$ or $k-i$ would be strictly larger than $j-i$, contradicting our choice of $v_{i}, v_{j}$ and $x$.

### 3.5 Cleanup of current vertex

Suppose we are extending the chordless path $\pi_{s u}$, while cleaning the st-path certificate. We identify a good vertex $v \in N(u)$, which closest to $t$ along the st-path. Ideally, we would clean the vertex $u$ by throwing away all its other neighbours but this could cost $\Omega(m)$ per chordless path as illustrated in Fig. 1 and discussed in Section 3.3. We thus perform a partial cleaning, called cleanup, which consists in identifying and removing, among all neighbours of $u($ i.e. $|N(u)|$ elements) only its set $N^{\text {good }}(u)$ of good ones.

For a given $u$ in a chordless st-path $\pi_{s t}=v_{0} \ldots v_{i} u \ldots v_{l}$, we let emerge the good neighbours in $N^{\text {good }}(u)$ one by one as follows. Consider the graph $G^{\prime}$ where the vertices $v_{0} \ldots v_{i}$ and its good neighbours were removed. If $u$ and $t$ are not connected, then there cannot be further chordless paths from $u$ and so there cannot be further good neighbours. Otherwise, if $u$ and $t$ are connected, we take any path from $u$ to $t$, and select its neighbour $v$ that appears along the path and is closest to $t$, as illustrated in Fig. 3. After that, we remove $v$ and its incident edges, and iterate what described above until $u$ is disconnected from $t$. The vertices $v$ thus selected form the set $N^{\text {good }}(u)$ of good neighbours.


Fig. 3. Cleanup of the neighbours of vertex $u$.

Lemma 2. For a chordless path $\pi_{s t}$, the cleanup of vertex $u \in \pi_{s t}$ correctly produces the set $N^{\text {good }}(u)$ of its good neighbours.

Proof. Given a chordless st-path $\pi_{s t}=v_{0} v_{i} u \ldots v_{\ell}$ in the graph $G$, for some integer $\ell>1$, where $v_{0}=s$ and $v_{\ell}=t$. Let $G^{\prime}$ be the subgraph of $G$ where the vertices $\left\{v_{0}, \ldots, v_{i}\right\}$ and their good neighbours were removed. Let $S \subseteq N\left(v_{i}\right)$ be the set of vertices removed by the cleanup procedure for $u$ in $G^{\prime}$. We divide the proof in two parts. We first show that $S$ contains $N^{\text {good }}\left(v_{i}\right)$, and then show that $N^{\text {good }}\left(v_{i}\right)$ contains $S$.

Clearly, for the first part, it is enough to show that $N_{G^{\prime}}\left(v_{i}\right) \backslash S$ are not good neighbours. The vertices of $N_{G^{\prime}}\left(v_{i}\right) \backslash S$ cannot reach $t$ in $G^{\prime}$ without passing through some vertex in $S$, otherwise $u$ would still be connected to $t$ and the procedure would not stop. This implies that there is no chordless ut-path in $G^{\prime}$ using some vertex of $N_{G^{\prime}}\left(v_{i}\right) \backslash S$, i.e. they are not good neighbours of $v_{i}$.

Finally, for the second part, it is enough to show that for each $w \in S$ there is a chordless $u t$-path in $G^{\prime}$ passing through $w$. Consider the iteration where $w \in N_{G^{\prime}}(u)$ was added to $S$ and let $S^{\prime} \subseteq S$ be the set of vertices added in previous iterations. We have that $w$ is the vertex closest to $t$ in the path $\pi=u \rightsquigarrow t$ in $G^{\prime} \backslash S^{\prime}$. We claim that any subpath of $(u, w) \pi_{w t}$ in $G^{\prime} \backslash S^{\prime}$ contains $(u, w)$, where $\pi_{w t}=w \rightsquigarrow t$ is a suffix of $\pi$. Thus implying that $w$ is contained in an induced $u t$-path in $G^{\prime}$. Indeed, the edge $(u, w)$ is not contained in a subpath iff there is a vertex $x \in N_{G^{\prime}}(u)$ in $\pi_{w t}$. By construction, $\pi_{w t}$ does not contain any vertex of $S^{\prime}$; and by the choice of $w, \pi_{w t}$ does not contain any vertex of $N_{G^{\prime}}(u) \backslash S^{\prime}$.

## 4 Listing Algorithm

We blend the key ideas discussed in Section 3 to get Algorithm [1] which has four parameters as input and lists all the chordless st-paths: the first parameter is the chordless path $\pi_{s u}$ partially built from $s$ to the current vertex $u$ (initially, $u=s$ ), which is the second parameter; the third parameter is a $u t$-path $\pi_{u t}$ that plays the role of certificate by Fact the fourth parameter is the reduced graph $G$, which changes with the recursive calls.

The algorithm outputs a chordless st-path if $u=t$ (line 24). Otherwise, it performs a cleanup of $u$ (the loop at lines 5-12). After that, it explores only the good neighbours recursively as they will surely lead to further chordless paths (the other loop at lines 13-17). Observe that $S$ stores the good neighbours $v$ of $u$ and a $v t$-path for each of them: when performing the recursive call at line 15, only

```
Algorithm 1: list_induced_paths \({ }_{s, t}\left(\pi_{s u}, u, \pi_{u t}, G\right)\)
    if \(u=t\) then
        output \(\left(\pi_{s u}\right)\)
    else
        \(S:=\emptyset\)
        while true do
            \(v:=\) the vertex in \(\pi_{u t} \cap N(u)\) that is closest to \(t\) in \(\pi_{u t}\)
            \(\pi_{v t}:=\) the subpath of \(\pi_{u t}\) from \(v\) to \(t\)
            \(S:=S \cup\left\{\left(v, \pi_{v t}\right)\right\}\)
            remove \(v\) and its incident edges from \(G\)
            if \(u\) and \(t\) are not connected then break
            \(\pi_{u t}:=\) any path from \(u\) to \(t\)
        end
        foreach \(\left(v, \pi_{v t}\right) \in S\) do
            adds back \(v\) and its incident edges to \(G\)
            list_induced_paths \({ }_{s, t}\left(\pi_{s u} \cdot(u, v), v, \pi_{v t}, G\right)\)
            remove \(v\) and its incident edges from \(G\)
        end
    end
```

one of the vertices in $S$ appears in the reduced graph $G$ passed as a parameter to the recursive call (see lines 14 and 16 that guarantee this, and Remark 1 below). Hence, the recursive call now has as parameters the chordless $s v$-path $\pi_{s u} \cdot(u, v)$ ending in $v$, and a $v t$-path that guarantees that a chordless $s t$-path exists and has $\pi_{s u} \cdot(u, v)$ as a prefix. This recursive call lists all the chordless st-paths that share this prefix.


Fig. 4. Two example graphs where $s=v_{0}$ and $t=v_{4}$.

For example, let us run Algorithm 1 on the input graph shown on the left of Fig. [4 with $u=s=v_{0}$ and the initial path $\pi_{u t}=v_{0} v_{1} v_{2} v_{3} v_{4}$. It computes the pairs $\left(v, \pi_{v t}\right)$ in $S$ as follows. First, $\left(v_{3}, v_{3} v_{4}\right)$ is added to $S$ as $v_{3}$ is a good neighbour for $v_{0}$ (the neighbour closest to $t$ in the path), and the edges incident to $v_{3}$ are removed. After this removal, $s=v_{0}$ is still connected to $t=v_{4}$ through the path $v_{0} v_{5} v_{6} v_{4}$, which becomes the input for the next iteration of the while
loop. Next, $\left(v_{6}, v_{6} v_{4}\right)$ is added to $S$ as $v_{6}$ is another good neighbour, and the edges incident to $v_{6}$ are removed disconnecting $s$ from $t$, so the while loop ends. The recursive calls in the foreach loop give the two chordless paths $v_{0} v_{3} v_{4}$ and $v_{0} v_{6} v_{4}$ contained in the graph.
Remark 1. It is important to run the recursive calls with all good neighbours in $S$ removed except one. If we left two or more good neighbours in the recursive call of line 15, they could interfere with each other and we might not obtain the chordless paths correctly. A very simple example is given in the graph shown on the right of Fig. [4. Consider for instance the case where Algorithm 1 would be given as input the path $v_{0} v_{1} v_{3} v_{2} v_{4}$. The pair $\left(v_{2}, v_{2} v_{4}\right)$ is added to $S$ and $v_{2}$ removed. After that, the path $v_{0} v_{1} v_{3} v_{5} v_{4}$ is found and the pair $\left(v_{1}, v_{1} v_{3} v_{5} v_{4}\right)$ is added to $S$ and $v_{1}$ removed. Since $v_{0}$ and $v_{4}$ become disconnected, $S$ contains all the good neighbours of $v_{0}$. Algorithm 1 executes the recursive calls with $S$. Suppose that we keep both good neighbours $v_{1}$ and $v_{2}$ in $G$ during these calls, in particular for the call with the pair $\left(v_{1}, v_{1} v_{3} v_{5} v_{4}\right)$ from $S$. This call will extend in a nested call the chordless path to $\pi_{s u}=v_{0} v_{1} v_{3}$ for $u=v_{3}$, and will claim that the good neighbours of $v_{3}$ are $v_{2}$ and $v_{5}$, which is incorrect since $v_{0} v_{1} v_{3} v_{2}$ is not chordless. This situation does not arise if $v_{2}$ is kept deleted in $G$ when the recursive call on $v_{1}$ is performed as done in Algorithm 1 .

The correctness of Algorithm 1 follows mostly from Lemma 2. Recall that, it guarantees that for a given path prefix $\pi_{s u}$ the set $S$ contains the good neighbours of $u$, i.e. the neighbours of $u$ that belong to at least one chordless st-path extending $\pi_{s u}$. Clearly, we only have to recursively call the algorithm for these neighbours, the others certainly lead to no solution. This implies that Algorithm 1 tries all the possibilities to extend $\pi_{s u}$, so all chordless $s t$-paths are output. Moreover, since each good neighbour of $u$ leads to a different extension, we have that no st-path is output more than once.

Certainly only st-paths are output by Algorithm 1 but at this point we have no guarantees that the paths are indeed chordless. In fact, after building $S$, the algorithm proceeds to recursively extend the prefix $\pi_{s u} \cdot(u, v)$ for each $v \in S$ in the graph $G^{\prime}=G \backslash(S \backslash\{v\})$. However, since $u$ was included in current path none of its neighbours can be used later in the recursion. The algorithm removes the good neighbours of $u$ from $G$, but the other neighbours, $N_{G}(u) \backslash S$, are still present in $G^{\prime}$. They could thus be used to extend the path later in the recursion, resulting in a non-chordless st-path. Lemma 3 shows that this cannot happen.
Lemma 3. The st-paths output by Algorithm 1 are chordless.
Proof. We proceed by contradiction. Suppose $\pi_{s t}$ is output by the algorithm and it is not chordless. This means that there exists a pair of vertices $x, y \in \pi_{s t}$ such that $x \neq y,(x, y)$ is an edge of $G$, and $(x, y) \notin \pi_{s t}$. We can assume the edge $(x, y)$ is such that $x$ is the vertex closest to $s$. Now, consider the recursive call corresponding to the prefix $\pi_{s x}$ : let $G_{x}$ be the associated graph and $z$ the next vertex of $\pi_{s t}$. The suffix $\pi_{x t}$ of $\pi_{s t}$ must pass through $z$, since $x$ is not connected to $t$ in $G_{x}-S$. So, $y$ is closest to $t$ than $z$ in $\pi_{s t}$. Thus, the path $(x, y) \pi_{y t}$, where $\pi_{y t}$ is a suffix of $\pi_{s t}$, avoids $S$ in $G_{x}$. This contradicts the test in line 10 .

The previous lemma leads to the following theorem.
Theorem 1. The algorithm correctly outputs all chordless st-paths of $G$.
Theorem 2. The algorithm takes $O\left(m+\left|\mathcal{P}_{s t}(G)\right|\left(t_{p}+n t_{q}+n t_{u}\right)\right)$ time, where $t_{p}$ is the cost of choosing any path from any two given vertices, $t_{q}$ is the cost of checking if any given two vertices are connected or not, and $t_{u}$ is the cost of removing/adding back any given edge.

Proof. See Section 5 .
There are several dynamic data structures in the literature 8 that maintain a spanning forest for a dynamic graph, supporting insertions and deletions of edges in polylogarithmic time. Consequently, $t_{p}=O(n \operatorname{polylog}(n))$, $t_{q}=O(\operatorname{polylog}(n))$, and $t_{u}=O($ polylog $(n))$, thus giving the following bound.

Corollary 1. The algorithm takes $\tilde{O}\left(m+\left|\mathcal{P}_{s t}(G)\right| \cdot n\right)$ time to report all the chordless st-paths.

## 5 Amortized Analysis

Before starting our analysis, we observe some simple properties of the recursion tree generated by Algorithm 1 .

Fact 2 The recursion tree $R$ of Algorithm 1 has the following properties:

1. There is a one-to-one correspondence between paths in $\mathcal{P}_{s t}(G)$ and leaves in the recursion tree.
2. There is a one-to-one correspondence between proper prefixes of paths in $\mathcal{P}_{s t}(G)$ and internal nodes in the recursion tree.
3. The number of branching nodes is $\left|\mathcal{P}_{s t}(G)\right|-1$.
4. The length of a root-to-leaf path is equal to the length of the chordless st-path corresponding to the leaf. In particular, the height of the tree is $\leq n$.

Fact 2 suggests us to follow the following overall strategy.

1. We analyze the cost of each type (leaf, unary and branching) of node separately.
2. We consider all branching nodes together, and show that their amortized cost is $O\left(t_{p}+t_{q}+n t_{u}+n\right)=\tilde{O}(n)$ per solution.
3. We consider all unary nodes together, and show that their amortized cost is $O\left(\left|\pi_{s t}\right| t_{q}+n t_{u}\right)=\tilde{O}(n)$ per solution.
4. We deduce that the cost of each solution is $O\left(t_{p}+n t_{q}+n t_{u}\right)=\tilde{O}(n)$.

Where the cost of a node is the time spent by the corresponding call without including the time spent by its nested recursive calls.

Lemma 4. The cost of a leaf is $O\left(\left|\pi_{s t}\right|\right)$.

Proof. Clearly, when $u=t$, the only operation done by the algorithm is to output $\pi_{s t}$, which takes $O\left(\left|\pi_{s t}\right|\right)$ time.

Let us now analyze the cost of the unary nodes. Let $r=\left\langle\pi_{s u}, u, \pi_{u t}, G\right\rangle$ be a unary node. The vertex $v \in N(u)$ is the only neighbour of $u$ that can extend the prefix $\pi_{s u}$ into a chordless st-path. Thus, removing $v$ from $G$ disconnects $u$ from $t$, and the algorithm performs a single iteration of the loop in line 5, not executing line 11, In this case, the algorithm performs the following operations: (i) one connectivity query (line 10), (ii) $|N(v)|$ edge update operations on $G$ (lines 9,14 and (16), and (iii) a scan in the intersection of $N(u)$ and $\pi_{u t}$ to find $v$ (line 6). The cost of (i) and (ii) is $O\left(t_{q}+|N(v)| t_{u}\right)$.

A naive implementation of (iii) takes $O\left(|N(u)|+\left|\pi_{u t}\right|\right)$ time, which is too large to fit in our amortization strategy. In order to reduce this cost to $O(|N(u)|)$ we therefore maintain, as an extra invariant, for each vertex in the current graph its distance to $t$ along the path $\pi_{u t}$. In this way, we can find $v$ simply scanning $N(u)$. Thus, assuming the distance information is correctly maintained, we complete the proof of Lemma 5 .

Lemma 5. The cost of a unary node is $O\left(t_{q}+|N(v)| t_{u}+|N(u)|\right)$, where $(u, v)$ is the edge added to the chordless path.

It is not hard to maintain the distance information for $\left\langle\pi_{s u} \cdot(u, v), v, \pi_{v t}, G^{\prime}\right\rangle$, the only child of the unary node $\left\langle\pi_{s u}, u, \pi_{u t}, G\right\rangle$. As the path $\pi_{v t}$ is a suffix of $\pi_{u t}$, the distance of the vertices in $\pi_{u t}$ does not change. On the other hand, the only vertices that the distances can change are the ones in $\pi_{v t}$ but not in $\pi_{u t}$. These vertices can be identified when scanning $N(v)$ in the child node $\left\langle\pi_{s u} \cdot(u, v), v, \pi_{v t}, G^{\prime}\right\rangle$, since their distance is strictly larger than $\left|\pi_{v t}\right|$. It remains to show that the distance information can be maintained in the branching nodes.

Lemma 6. The cost of a branching node $r \in R$ is $O\left(\beta(r)\left(t_{p}+t_{q}+n t_{u}\right)\right)$, where $\beta(r)$ is the number of children of $r$.

Proof. The cost of a branching node $r=\left\langle\pi_{s u}, u, \pi_{u t}, G\right\rangle$ is dominated by the cost of the loop of line 5. The number of iterations of the loop is equal to the number of neighbours of $u$ that can extend $\pi_{s u}$ into a chordless $s t$-path, which is exactly the number of vertices in $S$ after the loop finishes, i.e. $\beta(r)$, the number of children of $r$. Let us now bound the cost of each iteration. The cost of lines 6 and $\left[7\right.$ is bounded by $O\left(|N(u)|+\left|\pi_{u t}\right|\right)=O(n)$ : we simply have to traverse the path $\pi_{u t}$ and scan the set $N(u)$. The cost of updating $G$ is bounded by $O\left(n t_{u}\right)$. Finally, the cost for the connectivity query and to find a path is $t_{p}+t_{q}$. Hence, the total cost for a branching node is $O\left(\beta(r)\left(t_{p}+t_{q}+n t_{u}\right)\right)$.

Let us now show that we can maintain the distance information in branching nodes in the same time bound of Lemma 6. This follows from the fact that in each iteration of the loop (line 5) we are already paying $O\left(\left|\pi_{u t}\right|\right)$, i.e. a full traversal of the path $\pi_{u t}$. Before each recursive call in line 15 we can traverse the path $\pi_{u t}$ adding for each vertex the distance information, i.e. their position in the path.

At this point we have bounds for the cost of each node in the recursion tree. However, by directly applying them we cannot achieve our goal of $\tilde{O}(n)$ time per solution. For instance, consider the particular case where all internal nodes of the recursion tree are branching. The cost of each internal node is $O\left(\beta(r)\left(t_{p}+t_{q}+n t_{u}\right)\right)=\tilde{O}\left(n^{2}\right)$, since $\beta(r)=\Omega(n)$ in the worst case. Then, from item 3 of Fact 2, the number of branching nodes is $\left|\mathcal{P}_{s t}(G)\right|-1$. The total cost for the tree is thus $\tilde{O}\left(\left|\mathcal{P}_{s t}(G)\right| \cdot n^{2}\right)$ or $\tilde{O}\left(n^{2}\right)$ per solution.

In order to get a tighter bound for the total cost of the branching nodes, we use the following amortization strategy. Let $r \in R$ be a branching node. We divide the cost $O\left(\beta(r)\left(t_{p}+t_{q}+n t_{u}\right)\right)$ among the closest descendents that are branching nodes or leaves (no unary nodes), each being charged $O\left(t_{p}+t_{q}+n t_{u}\right)$. This can always be done since $r$ has $\beta(r)$ children and the subtree of each child contains at least one leaf, i.e. the node $r$ has at least $\beta(r)$ non-unary descendants. In this way, the original cost of node $r$ is completely charged to its non-unary descendants, and the only cost that remains associated to $r$ is the one received from its ancestors. Finally, each branching node can only be charged once, by its lowest non-unary ancestor. Each branching node and each leaf is therefore charged with $O\left(t_{p}+t_{q}+n t_{u}\right)$. Thus, the total cost of the branching nodes is $O\left(\left|\mathcal{P}_{s t}(G)\right|\left(t_{p}+t_{q}+n t_{u}\right)\right)$, completing the proof of Lemma 7 ,

Lemma 7. $\sum_{r: b r a n c h i n g} T(r)=O\left(\left|\mathcal{P}_{s t}(G)\right|\left(t_{p}+t_{q}+n t_{u}\right)\right)$.
Let us now bound the total cost of the unary nodes. Similarly to the branching nodes case, a straightforward use of the bound given by Lemma 5 leads to an $\tilde{O}\left(n^{2}\right)$ cost per solution, since in the worst case the recursion tree can have $O(n)$ unary nodes for each leaf. The key idea to obtain a better amortized cost is to consider the bound on the reduced degrees given by Lemma 1 .

We first observe that each unary node is contained in some root-to-leaf path $\Pi(l)$, where $l$ is a leaf of the recursion tree. Thus,

$$
\begin{equation*}
\sum_{r: \text { unary }} T(r) \leq \sum_{l: \text { leaf }} \sum_{r \in \Pi(l)} T(r) \tag{1}
\end{equation*}
$$

Fact 2 implies that there is a one-to-one correspondence between the prefixes of paths in $\mathcal{P}_{s t}(G)$ and nodes in the recursion tree. That is, each leaf corresponds to a solution, and the root-to-leaf path $\Pi(l)$ corresponds to the chordless stpath associated to the leaf $l$. Moreover, the $O\left(t_{q}+|N(v)| t_{u}+|N(u)|\right)$ cost of an unary node can be amortized to $O\left(t_{q}+|N(v)| t_{u}\right)$, since we can always charge $|N(u)|=O(n)$ to its single child. We can thus rewrite the double sum as

$$
\begin{equation*}
\sum_{l: \text { leaf }} \sum_{r \in \Pi(l)} T(r)=\sum_{\pi \in \mathcal{P}_{s t}(G)} \sum_{v_{i} \in \pi}\left(t_{q}+\left|N\left(v_{i}\right)\right| t_{u}\right) . \tag{2}
\end{equation*}
$$

For each chordless st-path $\pi$ in the internal sum of Eq. 2 we have that the degrees are actually the reduced degrees of Section 3.4. since the good neighbours (i.e. the set $S$ in Algorithm (1) are always removed. Using Lemma we can thus
bound the sum of the degrees by $2 n$. Therefore,

$$
\begin{equation*}
\sum_{r: \text { unary }} T(r) \leq \sum_{\pi \in \mathcal{P}_{s t}(G)}\left(|\pi| t_{q}+2 n t_{u}\right) \tag{3}
\end{equation*}
$$

completing the proof of Lemma 8 .
Lemma 8. $\sum_{r: \text { unary }} T(r)=O\left(\sum_{\pi \in \mathcal{P}_{s t}(G)}|\pi| t_{q}+n t_{u}\right)$.
As a corollary of Lemmas 8 and 7, we obtain Theorem 2.

## References

1. Etienne Birmelé, Rui A. Ferreira, Roberto Grossi, Andrea Marino, Nadia Pisanti, Romeo Rizzi, and Gustavo Sacomoto. Optimal listing of cycles and st-paths in undirected graphs. In SODA 2013, 1884-1896. ACM/SIAM, 2013.
2. Yijia Chen and Jörg Flum. On parameterized path and chordless path problems. In IEEE Conference on Computational Complexity, pages 250-263, 2007.
3. Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. Annals of Mathematics, 164:51-229, 2006.
4. Michele Conforti, Gérard Cornuéjols, Ajai Kapoor, and Kristina Vuskovic. Recognizing balanced $0,+/-$ matrices. In SODA 1994, 103-111. ACM/SIAM, 1994.
5. Michele Conforti, Gérard Cornuéjols, Ajai Kapoor, and Kristina Vuskovic. Finding an even hole in a graph. In FOCS 1997, 480-485. IEEE Computer Society, 1997.
6. Michele Conforti and M. R. Rao. Structural properties and decomposition of linear balanced matrices. Math. Program., 55:129-168, 1992.
7. Robert Haas and Michael Hoffmann. Chordless paths through three vertices. Theoretical Computer Science, 351(3):360-371, 2006.
8. Bruce M. Kapron, Valerie King, and Ben Mountjoy. Dynamic graph connectivity in polylogarithmic worst case time. In SODA, pages 1131-1142, 2013.
9. Ken-ichi Kawarabayashi and Yusuke Kobayashi. The induced disjoint paths problem. In Andrea Lodi, Alessandro Panconesi, and Giovanni Rinaldi, editors, IPCO, volume 5035 of Lecture Notes in Computer Science, pages 47-61. Springer, 2008.
10. R C Read and Robert E Tarjan. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. Networks, 5(3):237252, 1975.
11. D. Seinsche. On a property of the class of n-colorable graphs. Journal of Combinatorial Theory, Series B, 16(2):191-193, 1974.
12. Nayla Sokhn, Richard Baltensperger, Louis-Felix Bersier, Jean Hennebert, and Ulrich Ultes-Nitsche. Identification of chordless cycles in ecological networks. In COMPLEX, 2012.
13. Maciej M. Syslo. An efficient cycle vector space algorithm for listing all cycles of a planar graph. SIAM J. Comput., 10(4):797-808, 1981.
14. Takeaki Uno. Algorithms for enumerating all perfect, maximum and maximal matchings in bipartite graphs. In ISAAC 1997, LNCS, 92-101. Springer, 1997.
15. Takeaki Uno. An output linear time algorithm for enumerating chordless cycles. In 92nd SIGAL of Information Processing Society Japan, pages 47-53, 2003. (in Japanese).
16. Marcel Wild. Generating all cycles, chordless cycles, and hamiltonian cycles with the principle of exclusion. J. of Discrete Algorithms, 6(1):93-102, March 2008.

[^0]:    * GS and MFS were partially supported by the ERC programme FP7/2007-2013 / ERC grant agreement no. [247073]10, and the French project ANR-12-BS02-0008 (Colib'read).

