

Tracking Paths in Planar Graphs

David Eppstein

Department of Computer Science, University of California, Irvine, USA
eppstein@uci.edu

Michael T. Goodrich 

Department of Computer Science, University of California, Irvine, USA
goodrich@uci.edu

James A. Liu

Department of Computer Science, University of California, Irvine, USA
jamesall@uci.edu

Pedro Matias 

Department of Computer Science, University of California, Irvine, USA
pmatias@uci.edu

Abstract

We consider the NP-complete problem of tracking paths in a graph, first introduced by Banik et al. [4]. Given an undirected graph with a source s and a destination t , find the smallest subset of vertices whose intersection with any $s - t$ path results in a unique sequence. In this paper, we show that this problem remains NP-complete when the graph is planar and we give a 4-approximation algorithm in this setting. We also show, via Courcelle's theorem, that it can be solved in linear time for graphs of bounded-clique width, when its clique decomposition is given in advance.

2012 ACM Subject Classification Mathematics of computing → Graph theory; Theory of computation → Computational complexity and cryptography; Theory of computation → Design and analysis of algorithms

Keywords and phrases Approximation Algorithm, Courcelle's Theorem, Clique-Width, Planar, 3-SAT, Graph Algorithms, NP-Hardness

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2019.54

Acknowledgements We thank Nil Mamano for suggesting the problem of tracking paths on a graph.

1 Introduction

Motivated by applications in surveillance and monitoring, Banik et. al. [4, 3] introduced the problem of tracking paths in a graph. In essence, the goal is to uniquely determine the path traversed by a moving subject or object, based on a sequence of vertices sampled from that path. Examples of surveillance applications include the following: (i) vehicle tracking in road networks; (ii) habitat monitoring; (iii) intruder tracking and securing large infrastructures; (iv) tracing back of illicit Internet activities by tracking data packets. Another application would be to determine the nodes in a network that have been compromised by a spreading infection, given an incomplete transmission history of a pathogen. This information can be helpful in identifying and attenuating the negative impact caused by biological and non-biological infectious agents, such as:

- Highly-contagious diseases (e.g. Severe Acute Respiratory Syndrome) which could lead to epidemics [17].
- Fake news and hate speech being disseminated in social networks, as well as violations of privacy (e.g. sharing without permission highly sensitive content owned by a user, such as intimate pictures).
- Computer viruses, which spread throughout servers scattered across the Internet.



© David Eppstein, Michael T. Goodrich, James A. Liu, and Pedro Matias;
licensed under Creative Commons License CC-BY

30th International Symposium on Algorithms and Computation (ISAAC 2019).

Editors: Pinyan Lu and Guochuan Zhang; Article No. 54; pp. 54:1–54:17

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

54:2 Tracking Paths in Planar Graphs

Some of these applications have been studied empirically or using heuristics in [6, 18, 25, 28] or, in the case of infections, [22, 23, 27, 1]. To the best of our knowledge, Banik et. al. were the first to approach the problem of tracking paths from a theory perspective. In this work, we extend some of their work and give new algorithms. Our main results apply to graphs that can: be embedded on the plane (of interest to surveillance in road networks and similar infrastructures) or that have bounded clique-width (mainly of theoretical interest).

Preliminaries. In the tracking paths problem, we are given an undirected graph $G = (V, E)$ with no self loops or parallel edges and a source $s \in V$ and a destination $t \in V$. The goal is to place trackers on a subset of the vertices in a way that enables us to reconstruct exactly the path traversed from s to t . Let $T \subseteq V$ be a set of vertices (where we wish to place trackers) and let \mathcal{S}_P^T be the *sequence* of vertices in T visited during the traversal of a path¹ P . Let $u - v$ denote a path from u to v . We say that T is a *tracking set* if every $s - t$ path yields a unique sequence of observed vertices in T , that is, $\mathcal{S}_{P_1}^T \neq \mathcal{S}_{P_2}^T$ for all distinct $s - t$ paths P_1 and P_2 . We consider the following problem.

PLANAR-TRACKING (G, s, t)

Input: Undirected *planar* graph $G = (V, E)$ and two vertices $s \in V$ and $t \in V$.

Question: What is the smallest tracking set for G ?

We denote by TRACKING the problem of tracking paths when the input graph is not restricted to be planar. Due to space constraints, we defer proofs of Lemmas/Theorems marked with \star to the appendix.


Related work. Banik et. al. first introduced TRACKING in [3], where it is shown to be NP-hard by reducing from Vertex Cover, which seems unlikely to work in the planar case. Although not immediately obvious, they also show that TRACKING is in NP, by observing that every tracking set is also a *feedback vertex set*, i.e. a set of vertices whose removal yields an acyclic graph. Finally, they present a fixed-parameter tractable (FPT) algorithm (parameterized by the solution size) for the decision version of the problem, where they obtain a kernel of size $O(k^7)$ edges.

The concept of tracking set, however, first appeared in Banik et. al. [4], where they considered a variant of TRACKING that only concerns shortest $s - t$ paths, essentially modeling the input as a directed acyclic graph (DAG). Using a similar reduction from Vertex Cover, they show that this variant cannot be approximated within a factor of 1.3606, unless $P=NP$. They also give a 2-approximation for the planar version of tracking shortest paths, but they omit any hardness results for this variant.

More recently, Bilò et. al. [7] generalized the version of the problem concerning shortest $s - t$ paths, into the case of multiple source-destination pairs, for which they claim the first $O(\sqrt{n \log n})$ -approximation algorithm for general graphs. They also study a version of this multiple source-destination pairs problem in which the set of trackers itself (excluding the order in which they are visited) is enough to distinguish between $s - t$ shortest paths². In this setting, they claim a $O(\sqrt{n})$ -approximation algorithm and they show that it is NP-hard even for cubic planar graphs. The hardness construction intrinsically relies on the multiplicity of source-destination pairs and, therefore, cannot be adapted to the problem studied in this

¹ Some authors use the terms “path” and “walk” interchangeably, where vertices may be repeated, but in this paper, paths are required to have distinct vertices.

² Notice that when tracking shortest paths only and using a single source-destination pair, these two versions of the problem are the same.



■ **Figure 1** Entry-exit pair illustration, with entry vertex s' and exit vertex t' .

paper. They also give an FPT algorithm (w.r.t to the maximum number of vertices at the same distance from the source) for the problem concerning a single source-destination pair that was introduced in [4].

In [2], Banik and Choudhary generalize TRACKING into a problem on set systems³, which are characterized by a universe (e.g. vertex set) and a family of subsets of the universe (e.g. $s - t$ paths). They show that this generalized version of the problem is fixed-parameter tractable, by establishing a correspondence with the well known Test Cover problem.

Our results. In this paper, we give a 4-approximation for PLANAR-TRACKING (Section 3) and prove that it is NP-complete (Section 4). In addition, we show that TRACKING can be solved in cubic time for graphs of bounded clique-width and linear time if the clique decomposition of bounded width is given in advance (Section 5).

2 Definitions

► **Definition 1** (Entry-exit pair). *Let (G, s, t) be an instance of TRACKING. An entry-exit pair is, with respect to some simple cycle C in $G = (V, E)$, an ordered pair (s', t') of vertices in C that satisfy the following conditions:*

1. *There exists a path $s - s'$ from s to the entry vertex s'*
2. *There exists a path $t' - t$ from the exit vertex t' to t*
3. *Paths $s - s'$ and $t' - t$ are vertex-disjoint*
4. *Path $s - s'$ (resp. $t' - t$) and C share exactly one vertex: s' (resp. t').*

Essentially, an entry-exit pair (s', t') with respect to a cycle C (see Figure 1) represents two alternative $s - t$ paths and, thus, requires tracking at least one of them. We say that (s', t') is *tracked with respect to C* if and only if $C \setminus \{s', t'\}$ contains a tracker. In addition, C is *tracked* if and only if there is no entry-exit pair with respect to C that is untracked. If a cycle contains either (i) 3 trackers or (ii) s or t and 1 tracker in a non-entry/non-exit vertex, then it must be tracked. We say that these cycles are *trivially tracked*.

An alternative characterization of a tracking set, first given by Banik et al [3, Lemma 2], is the following.

► **Lemma 2** ([3]). *For a graph $G = (V, E)$, a subset $T \subseteq V$ is a tracking set if and only if every simple cycle C in G is tracked with respect to T .*

3 Approximation algorithm

► **Theorem 3.** *There exists a 4-approximation algorithm for PLANAR-TRACKING.*

The overall idea for the approximation algorithm builds on the following two insights:

1. The cardinality of the optimal solution cannot be much smaller than the number of faces in the graph.
2. The average number of trackers per face does not need to be very large.

³ Also called hypergraphs.

The first idea gives us a lower bound on OPT , the cardinality of an optimal solution. The second gives us an upper bound on ALG , the cardinality of our approximation algorithm.

3.1 Lower bound on OPT

We consider the following reduction, which takes care of disconnected components, or components that are “attached” to the graph by a cut vertex that is not in an $s - t$ path. We say that a reduction is *safe*, if it does not eliminate any untracked cycles.

Reduction 1. While there exists an edge or vertex that does not participate in any $s - t$ path, remove it from the graph.

► **Lemma 4** ([3]). ★ *Reduction 1 is safe and can be done in polynomial time.*

► **Lemma 5.** *After Reduction 1, every simple cycle in the graph contains at least one entry-exit pair. This holds for non-planar graphs as well.*

Proof. Let C be some simple cycle in the graph. After Reduction 1, there must exist an $s - t$ path that shares an edge with C . The first and last vertices on this path that belong to C correspond to an entry-exit pair. ◀

► **Lemma 6.** *In an embedded undirected planar graph G that results from Reduction 1, $OPT \geq (|F| - 1)/2$, where F is the set of faces of G .*

On a high level, the proof of Lemma 6 is done by keeping a set of “active” trackers while reconstructing a planar embedding \mathcal{E} of G : we start, as a base case, with any simple $s - t$ path in \mathcal{E} and iteratively add faces to it until it matches \mathcal{E} . Given a fixed, optimal tracking set T^* , the addition of each face requires either (i) adding a new tracker from T^* to the active set, or (ii) deactivating an active tracker, rendering it useless for distinguishing paths on future faces. As a consequence, each tracker charges at most two faces: the one adding the tracker and the one deactivating it. This demonstrates that $|T^*| \geq (|F| - 1)/2$.

Let $\text{OUTER}(\mathcal{E}^\tau)$ be the set of outer-edges of our planar reconstructing embedding \mathcal{E}^τ at time τ . At time $\tau = 0$, our embedding corresponds to an $s - t$ path and, for all $0 \leq \tau \leq |F| - 1$, we add exactly one face C from \mathcal{E} to \mathcal{E}^τ , by connecting two vertices u and v in $\text{OUTER}(\mathcal{E}^\tau)$ with a simple path p (see Figure 8 in the appendix). In doing so, we erase an $u - v$ path p' in $\text{OUTER}(\mathcal{E}^\tau)$, so we have that $\text{OUTER}(\mathcal{E}^{\tau+1}) = \text{OUTER}(\mathcal{E}^\tau) \setminus p' \cup p$. In the end, $\mathcal{E}^{|F|} = \mathcal{E}$.

By Lemma 5, there is at least one entry-exit pair in \mathcal{E} with respect to face C , so any tracking set must contain a tracker on some vertex of C . During the reconstruction process, we maintain a list of trackers in sets A and A' , such that $(A \cup A') \subseteq T^*$, where A contains active trackers and A' contains inactive ones. A tracker in vertex v is *active* at time τ if and only if it meets both of the following conditions:

Condition (i) $v \in \text{OUTER}(\mathcal{E}^\tau)$

Condition (ii) There is no $s - v$ path in \mathcal{E}^τ that traverses vertices in $\text{OUTER}(\mathcal{E}^\tau)$

Intuitively, an active tracker can be used to track future faces, although no more than one (see below). An inactive tracker, on the other hand, either cannot be used to track future faces (Condition (i)), or its corresponding vertex is entry/exit for some future face (Condition (ii)), in which case we require yet another tracker on that face (see Figure 2). Condition (ii) is necessary for dealing with embeddings of G where at least one of s and t is not in the outer face.



(a) Tracker on x (red) is inactive due to the violation of Condition (i).

(b) Tracker on x (red) is inactive due to the violation of Condition (ii). Notice that x is entry for exit vertices u and v (with respect to C), therefore C needs another tracker.

■ **Figure 2** Examples of inactive trackers used in the proof of Lemma 6.

Proof of Lemma 6.

First, we argue that each time we add a face during the reconstruction process described above, we either (i) need to increase the number of active trackers (by adding it to either A or A'), or (ii) we can get away by re-using and, therefore, deactivating an active tracker.

We assume for the rest of the argument that t is on the outer face, because such a planar embedding is always possible to construct.

Let C be the face added a time τ by connecting vertices u and v in $\text{OUTER}(\mathcal{E}^\tau)$, as specified above. In addition, let T^* be any optimal solution (i.e. $|T^*| = \text{OPT}$). We consider two cases, depending on the existence of a tracker in C at time τ :

Case 1: $C \cap A = \emptyset$ at time τ .

By Lemma 5, there exists a vertex $x \in C$ such that $x \in T^*$. We place a tracker on x . If $x \in \text{OUTER}(\mathcal{E}^{\tau+1})$ we add x to A , otherwise we add it to A' .

Case 1: $C \cap A \neq \emptyset$ at time τ .

Let $y \in C \cap A$ be a vertex of C with a tracker. We again consider two cases:

- (i) $y \notin \{u, v\}$. Then, $y \in \text{OUTER}(\mathcal{E}^\tau)$ but $y \notin \text{OUTER}(\mathcal{E}^{\tau+1})$, which amounts to moving y from A to A' .
- (ii) $y \in \{u, v\}$. If (u, v) is an entry-exit pair with respect to C , or if the tracker in y is not active, then there exists $x' \in C \setminus \{u, v\}$ such that $x' \in T^*$. Similarly to Case 1, we place a tracker on x' , which corresponds to adding x' either to A or A' .

Otherwise, the tracker in y is active and (u, v) is not an entry-exit pair with respect to C . Let us assume without loss of generality that $y = u$. Then, the addition of C deactivates the tracker in u by definition of active tracker (Condition (ii) is now violated), so we move u from A to A' .


Every tracker in A' is charged by at most two faces: one for adding an active tracker to A and another for deactivating it and moving it to A' . Therefore, $|F| - 1 \leq |A| + 2|A'|$. Since $|A| + |A'| \leq |T^*|$, it follows that $|F| - 1 \leq 2\text{OPT}$. ◀

A tight example for the lower bound on OPT is illustrated in Figure 3a.


3.2 Upper bound on ALG

We say that an undirected planar graph is *reduced* if it cannot be further reduced by Reduction 1 or any of the following reductions.

Reduction 2. While there exist two adjacent vertices of degree 2, remove one of them (and its edges) and add an edge connecting its neighbors.



(a) Example of a planar graph where $OPT = |F|/2$ (in red).



(b) Example of a planar graph where $ALG = 2(|F| - 2)$ (in red).

■ **Figure 3** Tight examples for the lower bound on OPT (left) and the upper bound on ALG (right) in planar graphs.

Reduction 3. While there exists vertex $v \notin \{s, t\}$ of degree 2 in a 3-cycle, place a tracker on v and remove it and its edges from the graph.

Reduction 4. While there exist non-adjacent vertices $u, v \notin \{s, t\}$ of degree 2 in a 4-cycle, place a tracker on either u or v and remove it and its edges from the graph.

Notice that all reduction rules are valid for general graphs, not only planar ones. In addition, Reductions 2, 3 and 4 can be applied interchangeably and in any order until none of them is applicable, but we will see that they need to be carried out after Reduction 1. Fortunately, we will not be required to re-apply Reduction 1 after performing the remaining reductions.


We denote the degree of a vertex v by $deg(v)$, where the underlying graph can be determined from its context.

▷ **Claim 7.** If vertex v is on an entry-exit pair, then $deg(v) > 2$.


Proof. Trivial by Definition 1. ◁

▷ **Claim 8.** Reductions 2, 3 and 4 maintain the property that every cycle in G contains at least one entry-exit pair (see Lemma 5).


Proof. Reductions 2, 3 and 4 only erase faces and vertices of degree 2, which cannot be in entry-exit pairs (by Claim 7), so every simple cycle of the graph still contains an entry-exit pair. ◁



(a) Illustration of Reduction 2, where $deg(u) = deg(v) = 2$.



(b) Illustration of Reduction 3, where $deg(s') \geq 3$, $deg(t') \geq 3$ and $deg(v) = 2$.



(c) Illustration of Reduction 4, where $deg(s') \geq 3$, $deg(t') \geq 3$ and $deg(u) = deg(v) = 2$.

■ **Figure 4** Illustration of Reductions 2, 3 and 4.

► **Lemma 9.** *★ Reduction 2 is safe and can be done in polynomial time, if done after Reduction 1.*

► **Lemma 10.** *★ Reduction 3 is safe and can be done in polynomial time, if done after Reduction 1.*

► **Lemma 11.** *★ Reduction 4 is safe and can be done in polynomial time, if done after Reduction 1.*

► **Remark 12.** None of Reductions 2, 3 and 4 compromise planarity.

■ **Algorithm 1** \mathcal{A} .

Input : Undirected planar graph $G = (V, E)$ and vertices $s \in V$ and $t \in V$

Output : Tracking set

- 1 Perform Reduction 1 in G
 - 2 Perform Reductions 2, 3 and 4 repeatedly until G is reduced.
 - 3 Output remaining vertices of degree at least 3 (except s or t)
-

► **Lemma 13.** *Algorithm \mathcal{A} outputs a tracking set for the input graph G .*

Proof. By Lemmas 4, 9, 10 and 11, Reductions 1-4 are safe, so let us assume without loss of generality that G is reduced. Then, every cycle of G of 5 or more vertices is trivially tracked, because it must contain at least 3 vertices of degree at least 3 (by Reduction 2). Similarly, every 3- or 4-cycle must contain at least 3 vertices of degree at least 3 by Reduction 2 and Reductions 3 and 4 (respectively). ◀

► **Lemma 14.** *Algorithm \mathcal{A} outputs a tracking set of size at most $2(|F| - 2)$, where F is the set of faces of the input graph G .*

Proof. Notice that each tracker added during Reductions 3 and 4 is associated with the removal of one face from G . Therefore, it is enough to show that the lemma holds with respect to a reduced graph G . Let us partition V into $V = (V_2 \cup V_{\geq 3})$, where V_2 and $V_{\geq 3}$ consist of the vertices of degree 2 and degree at least 3, respectively (notice that there cannot be vertices of degree 1). The lemma statement follows from Lemma 13 and the fact that $|F| \geq \frac{|V_{\geq 3}|}{2} + 2$. This inequality can be derived by plugging in the inequality $2|E| \geq 3|V_{\geq 3}| + 2|V_2|$ in Euler's formula for planar graphs: $|V| - |E| + |F| = 2$, where E is the set of edges of G . ◀


A tight example is illustrated in Figure 3b.

Proof of Theorem 3. By Lemmas 6 and 14, Algorithm \mathcal{A} is a 4-approximation to PLANAR-TRACKING. ◀

4 Hardness of tracking paths

We show that PLANAR-TRACKING is NP-hard, by reducing from PLANAR-3-SAT, a special version of the satisfiability problem, shown to be NP-complete by Lichtenstein [21].

In 3-SAT, we are given a set $\mathcal{X} = \{x_1, \dots, x_p\}$ of variables and a 3-CNF formula ϕ , where each clause in ϕ is a disjunction of exactly three distinct literals with respect to \mathcal{X} . The goal is to find a boolean assignment to all variables in \mathcal{X} that satisfies ϕ . Consider the bipartite



■ **Figure 5** Illustration of x_i 's gadget, containing m_i vertices for x_i (in blue) and m_i vertices for \bar{x}_i (in red). Vertices colored black require trackers in any minimum tracking set. The dashed edges are added to force trackers in s_i and t_i .

graph with a vertex for each clause C in ϕ and each variable $x_i \in \mathcal{X}$, and edges (x_i, C) if and only if C contains x_i or its negation \bar{x}_i . Lichtenstein [21] showed that PLANAR-3-SAT, the subset of instances of 3-SAT whose underlying bipartite graph is planar, remains NP-complete. In particular, the definition of PLANAR-3-SAT requires that a cycle can be drawn connecting all of the variables while maintaining planarity. Later, Knuth and Raghunatan [20] exploited this condition to show that we can always draw the underlying bipartite graph of a PLANAR-3-SAT instance in a *rectilinear* fashion without crossings (example in Figure 10a): variables are arranged in a horizontal line and clauses are horizontal line segments with vertical legs to represent the literals present in the clause. Vertical legs attach to the appropriate variables and are labeled *red* for negated literals and *blue*, otherwise. In particular, a given clause is drawn completely above or below the line of variables.

We convert a planar rectilinear drawing \mathcal{D} of an instance of PLANAR-3-SAT, with formula ϕ and a set \mathcal{X} of variables, into a planar drawing \mathcal{G} corresponding to the instance of PLANAR-TRACKING. The reduction is straightforward:

- (i) transform each variable x_i in \mathcal{D} into a gadget containing m_i copies of literal vertices x_i and \bar{x}_i ;
- (ii) transform each 3-legged clause into a face containing corresponding literals vertices and an entry-exit pair;
- (iii) choose the boolean assignment according to the placement of trackers, such that a clause is satisfied if and only if its corresponding face is tracked.

The union of all the variable and clause gadgets constitutes \mathcal{G} (see example in Figure 10 in the appendix). Details of each gadget are given below. For simplicity, we avoid introducing too many subscripts and we rely on pictures to describe the gadgets.

Variable gadget. Each variable gadget converts a variable x_i in \mathcal{D} into a connected subgraph corresponding to Figure 5, with length parameterized by m_i . We refer to the set $\{h_k, \mu_k, l_k\}$ as column k and we refer to the vertices $\{h_1, \dots, h_{m_i}\} \cup \{l_1, \dots, l_{m_i}\}$ as *literal vertices*.

Each variable gadget is linked with the next one by setting $t_i = s_{i+1}$, to form a horizontal chain of gadgets, where $s = s_1$ and $t = t_p$. For convenience, we force trackers in all the s_i (except s), by drawing edges between the α' (β') of a variable gadget and the α (β) of the next variable gadget in the chain. We refer to the resulting drawing as the *spine*.

There are exactly two minimum tracking sets associated with the variable gadget with source s_i and destination t_i . One of them corresponds to a true assignment of x_i and the other one to a false assignment. Both of them require tracking the vertices in $R =$

$\{\alpha, \alpha', \beta, \beta', \mu_1, \mu_{m_i}\}$, as well as the remaining μ_k . In addition, the true assignment tracks the even-indexed h_k and odd-indexed l_k , while the false assignment tracks the odd-indexed h_k and even-indexed l_k . This requires $2m_i + 4$ trackers in total.

► **Lemma 15.** *★ The true and false assignments are the only minimum tracking sets.*

Clause gadget. Let $C = (\ell_a \vee \ell_b \vee \ell_c)$ be a clause in ϕ with literals corresponding to variables $x_a, x_b, x_c \in \mathcal{X}$. Its gadget, depicted in Figure 6, is a face F_C consisting of:

- literal vertex α from x_a 's gadget, corresponding to literal ℓ_a ;
- adjacent literal vertices $\beta_1, \overline{\beta_1}, \beta_2, \overline{\beta_2}, \beta_3$ from x_b 's gadget, corresponding to literals alternating between ℓ_b and $\overline{\ell_b}$;
- literal vertex γ from x_c 's gadget, corresponding to literal ℓ_c ;
- edges $(\alpha, \gamma), (\alpha, \beta_1), (\beta_3, \gamma)$ and the edges from x_b 's gadget connecting all of the β_k and $\overline{\beta_k}$.

We can increase the lengths of the variable gadgets to any polynomial that provides enough literal vertices for all clauses. Since \mathcal{D} is planar, there are no crossings between clauses. We also impose the following restrictions:


1. α cannot be one of $\{h_1, l_{m_a}\}$; this ensures that the only faces in x_a 's gadget that do not require 3 trackers do not become untracked. We apply the equivalent restriction to β_1, β_3 and γ .
2. The $\alpha'_k/\overline{\alpha'_k}$, (see Figure 6) cannot belong to any other clause gadget; these correspond to the 4 literal vertices following α in x_a 's gadget and reserving them ensures that non-clause faces, between nested clauses, are tracked. We apply the equivalent restriction to the $\gamma'_k/\overline{\gamma'_k}$.
3. All literal vertices in a clause need to be on the same side of the spine; this restriction is trivial because \mathcal{D} is rectilinear, but it simplifies the analysis.

► **Lemma 16.** *★ Clause C is satisfied if and only if its corresponding gadget face F_C is tracked.*

► **Theorem 17.** *★ There exists a polynomial time reduction from PLANAR-3-SAT to PLANAR-TRACKING.*

► **Corollary 18.** PLANAR-TRACKING is NP-hard.

It remains to show that PLANAR-TRACKING is in NP; Banik et. al. [3] prove this in the more general case of TRACKING.



■ **Figure 6** Illustration of the gadget for clause $C = (\ell_a \vee \ell_b \vee \ell_c)$, where the vertices in each variable gadget are all adjacent. The entry-exit (β_1, β_2) are responsible for satisfying C .

5 Bounded clique-width graphs

We show that TRACKING can be solved in linear time when the input graph has bounded clique-width, by applying Courcelle’s theorem [9, 12, 14], a powerful meta-theorem that establishes fixed-parameter tractability of *any* graph property that is expressible in monadic second order logic.

Clique-width, first introduced by Courcelle et. al. [13] and revisited by Courcelle and Olariu [16], is an important graph parameter that, intuitively, measures the closeness of a graph to a cograph – a graph with no induced 4-vertex paths. It is closely related to *tree-width*, another influential graph parameter that measures closeness of a graph to a tree and that was first introduced by Bertelé and Brioschi [5] and later rediscovered by Halin [19] and Robertson and Seymour [26]. While both parameters are determined based on specific hierarchical decompositions of a graph, the clique-width is strictly more powerful in the sense that the class of graphs of bounded clique-width includes all graphs of bounded tree-width, but not vice-versa. Details on the relationship between these parameters can be found in Courcelle and Engelfriet [12]. Graphs of bounded clique-width include series-parallel graphs, outerplanar graphs, pseudoforests, cographs, distance-hereditary graphs, etc.


MSO₁ vs MSO₂. Second order logic extends first order logic, by allowing quantification over relations (of any fixed arity) on the elements of the domain of discourse. Monadic second order logic itself only allows quantification over unary relations (subsets of the domain of discourse) and, in the logic of graphs, it comes in two flavors: MSO₁ and MSO₂. The only distinction between these is that the latter allows edges to be elements of the domain of discourse (and thus be quantified over), while the former does not. Besides the quantifiers (\forall and \exists) and the standard logic operations $\neg, \wedge, \vee, \rightarrow$, both logics include predicates for equality ($=$) and relation membership (\in). In addition, MSO₁ includes a predicate (\sim) that determines vertex adjacency and MSO₂ includes a predicate for vertex-edge incidence. MSO₂ is more expressive, for example: Hamiltonicity can be expressed using MSO₂, but not using MSO₁. Details on the distinction between the two logics can be found in [12].

Courcelle’s theorem. Courcelle et. al. [9, 14] showed that any graph property expressed in MSO₁ and MSO₂ is FPT under clique-width and tree-width (respectively). More specifically, they showed that any MSO₁-, MSO₂-expressible property can be tested in $f(k, l) \cdot n$ and $g(k', l') \cdot (n + m)$ time (respectively) for graphs of clique-width k and tree-width k' , where: f and g are computable functions, l and l' are the lengths of the logic formulas, n is the number of vertices and m is the number of edges. The result for MSO₂ is valid in optimization problems with linear evaluation functions [15]. Later, Courcelle, Makowsky and Rotics [14] extended these results for MSO₁. Examples of constructing FPT graph algorithms parameterized by clique-width or tree-width, which are based on automata, are given in [10, 11]. While it is possible to construct tree decompositions of width k' in linear time [8], there is no FPT algorithm for finding clique decompositions of clique-width $k > 3$. Fortunately, it is possible to construct a clique decomposition of width exponential in k in cubic time [24]. In this section, we will take advantage of the latter.

We give an alternative definition for tracking set that is easier to express using the logic of graphs.

► **Lemma 19 (Tracking set).** *For an undirected graph $G = (V, E)$, a subset $T \subseteq V$ is a tracking set if and only if there is no $s - t$ path $P_{st} = P_{ss'} \cup P_{s't'} \cup P_{t't}$ for $s', t' \in V$ and corresponding $s - s'$, $s' - t'$ and $t' - t$ paths, such that:*

1. *There exists an alternative $s' - t'$ path $P'_{s't'} \neq P_{s't'}$ and*
2. *$T \cap (P_{s't'} \cup P'_{s't'}) \subseteq \{s', t'\}$.*



■ **Figure 7** Illustration of variable sets P , Q , Q' and R as well as vertex variables s , s' , t' and t used in expressing ISTRACKINGSET using MSO_1 .

Proof. This follows directly from Lemma 2. ◀

In the logic formulas presented below, we use lowercase letters to quantify over vertices and uppercase letters to quantify over sets of vertices. We use s and t as free variables and, for convenience, we also use set intersection (\cap), union (\cup) and containment (\subseteq), without explicitly expressing these operations using MSO_1 .

$$\begin{aligned} & \text{ISTRACKINGSET}(T, s, t) \\ \iff & \exists P, Q, R [\exists s', t' [\text{HASPATH}(P, s, s') \wedge \text{HASPATH}(Q, s', t') \wedge \text{HASPATH}(R, t', t) \\ & \wedge P \cap Q = \{s'\} \wedge Q \cap R = \{t'\} \wedge P \cap R = \emptyset \\ & \wedge \exists Q' \neq Q [\text{HASPATH}(Q', s', t') \wedge T \cap (Q \cup Q') \subseteq \{s', t'\}]]] \end{aligned}$$

The first two lines of the above equivalence establish that P , Q and R form an $s - t$ path. The last line restricts s' and t' to be an entry-exit pair with respect to the cycle $Q \cup Q'$ (see Figure 7) and, in addition, establishes that the cycle $Q \cup Q'$ is not tracked.

The primitive $\text{HASPATH}(X, a, b)$, whose input consists of a set $X \subseteq V$ and vertices $a, b \in V$, verifies the existence of a simple path between a and b that only uses vertices in X . We define it as follows:

$$\begin{aligned} & \text{HASPATH}(X, a, b) \\ \iff & \exists X_1, X_2 \subseteq X [X_1 \cup X_2 = X \wedge a \in X_1 \wedge b \in X_2 \wedge \neg(\exists u \in X_1 \wedge \exists v \in X_2 [u \sim v])] \end{aligned}$$

► **Remark 20.** $\text{HASPATH}(X, a, b)$ is correctly expressed under MSO_1 and it correctly verifies that there exists an $a - b$ path using only vertices in X .

► **Remark 21.** ISTRACKINGSET is correctly expressed under MSO_1 and it correctly verifies that the given subset of vertices is a tracking set.

► **Theorem 22.** $\text{TRACKING}(G, s, t)$ can be solved in polynomial time if G has bounded clique-width. Moreover, if a clique decomposition of bounded width is given, it can be solved in linear time.

Proof. This follows directly from Lemma 19, Remarks 20 and 21, and the linear time algorithm given by Courcelle [14] for any optimization problem on graphs of bounded clique-width, whose decomposition is given in advance. ◀

6 Conclusion and open questions

We showed that PLANAR-TRACKING is NP-complete and we give a 4-approximation algorithm. We also show that, for graphs of bounded-clique width, TRACKING can be solved in linear time by applying Courcelle's theorem, as long as its clique decomposition is given in advance. A natural direction of future study would be to improve the approximation ratio of PLANAR-TRACKING or establish constant approximation factors for graphs of larger genus or, more


generally, for TRACKING. Another open question is to establish the difficulty of TRACKING on directed graphs: on one side planar directed acyclic graphs are not known to be NP-hard and, on the other side, it is not known whether its general or planar versions are in NP. Finally, it would be interesting to find efficient algorithms for graphs of bounded tree-width or clique-width without resorting to the finite automaton approach used in Courcelle's theorem.

References

- 1 Norman TJ Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Charles Griffin & Company Ltd, High Wycombe, United Kingdom, 2nd edition, 1975.
- 2 Aritra Banik and Pratibha Choudhary. Fixed-Parameter Tractable Algorithms for Tracking Set Problems. In B. S. Panda and Partha P. Goswami, editors, *Algorithms and Discrete Applied Mathematics - 4th International Conference, CALDAM 2018, Guwahati, India, February 15-17, 2018, Proceedings*, volume 10743 of *Lecture Notes in Computer Science*, pages 93–104. Springer, 2018. doi:10.1007/978-3-319-74180-2_8.
- 3 Aritra Banik, Pratibha Choudhary, Daniel Lokshtanov, Venkatesh Raman, and Saket Saurabh. A polynomial sized kernel for tracking paths problem. In *Latin American Symposium on Theoretical Informatics*, volume 10807 of *Lecture Notes in Computer Science*, pages 94–107. Springer, 2018. doi:10.1007/978-3-319-77404-6_8.
- 4 Aritra Banik, Matthew J Katz, Eli Packer, and Marina Simakov. Tracking paths. In *International Conference on Algorithms and Complexity*, volume 10236 of *Lecture Notes in Computer Science*, pages 67–79. Springer, 2017. doi:10.1007/978-3-319-57586-5_7.
- 5 Umberto Bertele and Francesco Brioschi. *Nonserial Dynamic Programming*. Academic Press, 1972. URL: <https://www.elsevier.com/books/nonserial-dynamic-programming/bertele/978-0-12-093450-8>.
- 6 Sania Bhatti and Jie Xu. Survey of target tracking protocols using wireless sensor network. In *2009 Fifth International Conference on Wireless and Mobile Communications*, pages 110–115. IEEE, 2009. doi:10.1109/ICWMC.2009.25.
- 7 Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Tracking Routes in Communication Networks. In Keren Censor-Hillel and Michele Flammini, editors, *Structural Information and Communication Complexity - 26th International Colloquium, SIROCCO 2019, L'Aquila, Italy, July 1-4, 2019, Proceedings*, volume 11639 of *Lecture Notes in Computer Science*, pages 81–93. Springer, 2019. doi:10.1007/978-3-030-24922-9_6.
- 8 Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing*, 25(6):1305–1317, 1996. doi:10.1137/S0097539793251219.
- 9 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.
- 10 Bruno Courcelle and Irène Durand. Automata for the verification of monadic second-order graph properties. *J. Applied Logic*, 10(4):368–409, 2012. doi:10.1016/j.jal.2011.07.001.
- 11 Bruno Courcelle and Irène Durand. Computations by fly-automata beyond monadic second-order logic. *Theor. Comput. Sci.*, 619:32–67, 2016. doi:10.1016/j.tcs.2015.12.026.
- 12 Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic - A Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*. Cambridge University Press, 2012. URL: http://www.cambridge.org/fr/knowledge/isbn/item5758776/?site_locale=fr_FR.
- 13 Bruno Courcelle, Joost Engelfriet, and Grzegorz Rozenberg. Handle-rewriting hypergraph grammars. *Journal of Computer and System Sciences*, 46(2):218–270, 1993. doi:10.1016/0022-0000(93)90004-G.
- 14 Bruno Courcelle, Johann A Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000. doi:10.1007/s002249910009.


- 15 Bruno Courcelle and Mohamed Mosbah. Monadic second-order evaluations on tree-decomposable graphs. *Theoretical Computer Science*, 109(1-2):49–82, 1993. doi:10.1016/0304-3975(93)90064-Z.
- 16 Bruno Courcelle and Stephan Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(1-3):77–114, 2000. doi:10.1016/S0166-218X(99)00184-5.
- 17 Kee-Tai Goh, Jeffery Cutter, Bee-Hoon Heng, Stefan Ma, Benjamin KW Koh, Cynthia Kwok, Cheong-Mui Toh, and Suok-Kai Chew. Epidemiology and control of SARS in Singapore. *Annals of the Academy of Medicine, Singapore*, 35(5):301, 2006. PMID:16829997.
- 18 Rahul Gupta and Samir R Das. Tracking moving targets in a smart sensor network. In *2003 IEEE 58th Vehicular Technology Conference. VTC 2003-Fall*, volume 5, pages 3035–3039. IEEE, 2003. doi:10.1109/VETEFC.2003.1286181.
- 19 Rudolf Halin. S-functions for graphs. *Journal of Geometry*, 8(1-2):171–186, 1976. doi:10.1007/BF01917434.
- 20 Donald E Knuth and Arvind Raghunathan. The problem of compatible representatives. *SIAM Journal on Discrete Mathematics*, 5(3):422–427, 1992. doi:10.1137/0405033.
- 21 David Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982. doi:10.1137/0211025.
- 22 Cristopher Moore and Mark EJ Newman. Epidemics and percolation in small-world networks. *Physical Review E*, 61(5):5678, 2000. doi:10.1103/PhysRevE.61.5678.
- 23 Mark EJ Newman. Spread of epidemic disease on networks. *Physical Review E*, 66(1):016128, 2002. doi:10.1103/PhysRevE.66.016128.
- 24 Sang-Il Oum. Approximating rank-width and clique-width quickly. *ACM Transactions on Algorithms*, 5(1):10, 2008. doi:10.1145/1435375.1435385.
- 25 Tao Peng, Christopher Leckie, and Kotagiri Ramamohanarao. Survey of network-based defense mechanisms countering the DoS and DDoS problems. *ACM Computing Surveys*, 39(1):3, 2007. doi:10.1145/1216370.1216373.
- 26 Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, 1986. doi:10.1016/0196-6774(86)90023-4.
- 27 Devavrat Shah and Tauhid Zaman. Rumors in a network: Who’s the culprit? *IEEE Transactions on Information Theory*, 57(8):5163–5181, 2011. doi:10.1109/TIT.2011.2158885.
- 28 Alex C Snoeren, Craig Partridge, Luis A Sanchez, Christine E Jones, Fabrice Tchakountio, Stephen T Kent, and W Timothy Strayer. Hash-based IP traceback. *ACM SIGCOMM Computer Communication Review*, 31(4):3–14, 2001. doi:10.1145/383059.383060.

A Deferred proofs on the approximation algorithm



■ **Figure 8** Illustration of valid and invalid choices of vertices u and v in the proof of Lemma 6. Cycle C must correspond to a face in the fixed embedding \mathcal{E} .

54:14 Tracking Paths in Planar Graphs



■ **Figure 9** Illustration of Reductions 2, 3 and 4.

► **Lemma 4** ([3]). *★ Reduction 1 is safe and can be done in polynomial time.*

Proof. See proof in [3]. ◀

► **Lemma 9.** *★ Reduction 2 is safe and can be done in polynomial time, if done after Reduction 1.*

Proof. Banik et. al. [3, Lemma 16] showed that having three adjacent vertices of degree 2 is redundant. Here, we extend their proof and show that we can also get rid of the second adjacent vertex of degree 2. Let u and v be the two adjacent vertices of degree 2, and let u' be the other neighbor of u . Notice that v, u' are not adjacent, otherwise the u, v edge would then not belong to an $s - t$ path and thus be removed by Reduction 1. Hence, no parallel edges are added to the graph after this reduction. We argue that this reduction is safe by showing that there exists a minimum tracking set that does not include u and v simultaneously. Take any minimum tracking set T that includes u and v . We can always move the tracker from u to u' ; this remains a tracking set, because u immediately follows or precedes u' on any $s - t$ path. This can only decrease T 's cardinality. This reduction can be done in $O(\text{poly}(n))$ time, by repeatedly checking for the existence of adjacent vertices of degree 2. ◀

► **Lemma 10.** *★ Reduction 3 is safe and can be done in polynomial time, if done after Reduction 1.*

Proof. Let $v \notin \{s, t\}$ be the vertex of degree 2 in the triangle $\Delta vs't'$ (see Figure 9b). Then, it must be the case that s' and t' form an entry-exit pair. This follows from (i) the fact that there exists an entry-exit pair in $\Delta vs't'$ (by Lemma 5) and (ii) the fact that v cannot be in an entry-exit pair (by Claim 7). Then, by Lemma 2, any feasible solution must place a tracker on v . Therefore, v and its edges can be removed, since v is neither a cut-vertex, nor in any entry-exit pair, so that its removal does not eliminate an untracked cycle. Clearly, this reduction can be done in $O(\text{poly}(n))$ time. ◀

► **Lemma 11.** *★ Reduction 4 is safe and can be done in polynomial time, if done after Reduction 1.*

Proof. Let $u, v \notin \{s, t\}$ be the two vertices of degree 2 that connect the same pair of vertices s' and t' . Similarly to the proof of Lemma 10, s' and t' must form an entry-exit pair. So, by Lemma 2, any feasible solution must place a tracker in either u or v . By symmetry, we can track and remove u and its edges. Therefore, Reduction 4 is safe by Claim 7 and the facts that u is neither a cut-vertex nor in an entry-exit pair. ◀

B

 Deferred proofs on NP-completeness

To prove Lemma 15, we first prove a couple of helpful lemmas and make a few observations:

- Each vertex in R belongs to a triangle, such that the other two vertices form an entry-exit pair, so must be tracked.
- Each square face, besides the two including h_2 and l_{m_i-1} , requires 3 tracked vertices.
- Two adjacent vertices cannot both be untracked.

► **Lemma 23.** *The true/false assignments of x_i correspond to tracking sets with respect to x_i 's gadget with source s_i and destination t_i .*

Proof. In a true/false assignment of x_i , every face of the gadget is trivially tracked, except the faces including s_i/t_i (which are clearly tracked) and the faces including h_2 and l_{m_i-1} . Though the latter faces may only contain 2 trackers (depending on x_i 's truth value and/or m_i 's parity), they are nevertheless tracked because one of these trackers cannot be in an entry-exit pair. The remaining cycles, i.e. the ones which are not faces, are also trivially tracked: since these cycles must have size at least 6, they must contain 3 trackers given the observation that no adjacent vertices are untracked. ◀

► **Lemma 24.** *In a minimum tracking set, each column has exactly 2 trackers.*

Proof. Since the true/false assignments achieve this property, we only have to show that each column requires at least two trackers. Assume that a minimum tracking set has only one tracker in column $1 < k < m_i$; it must be on μ_k . Then all vertices on columns $k-1, k+1$ must be tracked. For the average number of trackers per column to be at most 2, the number of trackers per column must be an alternating sequence of $\dots 3, 1, 3, 1, \dots$, with column 1 and/or column m_i only having 1 tracker, which contradicts the square face property. ◀

► **Lemma 15.** *★ The true and false assignments are the only minimum tracking sets.*

Proof. Assume that some μ_k is not tracked in a minimum tracking set, for $1 < k < m_i$. By Lemma 24, we only have to show that this causes a column to have 3 trackers. If $k = m_i - 1$ then column m_i must have three trackers. Otherwise, $\mu_{k+1}, h_{k+1}, h_{k+2}$ must be tracked, since they share a square face with μ_k . If l_{k+1} is tracked we are done, otherwise, l_{k+2} must be tracked. Additionally, μ_{k+2} must be tracked, either by the square face property, or in the case where $k = m_i - 2$, because it is in R . Then, column $k+2$ has 3 trackers.

Now, if all the μ_k are tracked, then the only minimum tracking sets are the true and false assignments, by the observation that two adjacent vertices cannot both be untracked and Lemma 23. ◀

► **Lemma 16.** *★ Clause C is satisfied if and only if its corresponding gadget face F_C is tracked.*

Proof. The entry-exit $(\overline{\beta_1}, \overline{\beta_2})$ is the only one, with respect to F_C , that is tracked if and only if C is satisfied. The remaining entry-exit pairs are tracked by either a β_k or a $\overline{\beta_k}$. ◀

► **Theorem 17.** *★ There exists a polynomial time reduction from PLANAR-3-SAT to PLANAR-TRACKING.*

Proof. Let (G, \mathcal{G}) be an instance of PLANAR-TRACKING that results from applying the transformation described above to an instance (ϕ, \mathcal{D}) of PLANAR-3-SAT, where \mathcal{G} and \mathcal{D} are the underlying planar embeddings. We show that ϕ is satisfiable if and only if G has a tracking set of size $T = (\sum_{i=1}^p 2m_i + 4) + p - 1$.

- (\Leftarrow). Choose the truth assignment of each variable according to the given tracking set. The implication follows from Lemma 15 and the fact that if some clause in ϕ was not satisfied, then its gadget face would have been untracked, a contradiction.
- (\Rightarrow). Place T trackers on the literal vertices that correspond to the satisfiable truth assignment of every variable and on all non-literal vertices (except s and t). We show that this corresponds to a tracking set by arguing that every cycle C in G is tracked. We distinguish between two cases:


Case 1: C contains no clause edges.

Then C is tracked by (almost) the same argument given in Lemma 23 that shows that a truth assignment of x_i corresponds to a tracking set with respect to x_i 's gadget. Notice that, because of Restriction 1, clause edges are only added to faces that require 3 trackers, so this does not change the argument for the faces which do not require 3 trackers. The only differences are: (i) the addition of the edges that force trackers in every s_i , which only helps the argument, and (ii) the fact that C may span multiple variable gadgets, in which case C must traverse at least 3 trackers on the non-literal vertices between two variable gadgets.


Case 2: C contains clause edges.

Notice that, by construction of G , C must have at least one spine edge. If C corresponds to a clause face, then it must be tracked by Lemma 16. Otherwise, we show that C contains at least 3 trackers and, thus, is trivially tracked. Let us think of C as alternating non-empty paths of two types: *clause paths*, which only contain clause edges and *spine paths*, which only contain spine edges. To avoid dealing with complex cycles, we observe that each spine path in C must contain at least 1 tracker; this follows from the fact at least one of the 2 vertices sharing a spine edge must have a tracker (see variable gadget). Thus, let us assume that C contains no more than 2 spine/clause paths, or otherwise C immediately contains 3 trackers. Notice that if one of the spine paths spans 2 or more variable gadgets, then it must traverse at least 3 trackers on the non-literal vertices between two variable gadgets. Since every clause in ϕ contains exactly 3 distinct literals and C is simple, the only cases where none of the spine paths span multiple variable gadgets are the following:

- (i) C contains exactly 2 clause paths in different sides of the spine.
Then, the 2 spine paths connecting the two sides of the spine must traverse 2 trackers each (see variable gadget). Therefore, C contains at least 4 trackers.
- (ii) C contains exactly 2 clause paths on the same side of the spine, which both start or both end at the same variable gadget, one nested in the other.
Then, by Restriction 2 one of the spine paths contains at least 6 vertices, half of which must be tracked. \blacktriangleleft



(a) Example of a PLANAR-3-SAT instance drawn in a rectilinear fashion with clauses $A = (\overline{x_1} \vee x_2 \vee \overline{x_3})$, $B = (x_1 \vee \overline{x_4} \vee x_5)$, $C = (x_2 \vee \overline{x_3} \vee x_4)$, $D = (x_1 \vee \overline{x_2} \vee \overline{x_5})$.



(b) Example of a PLANAR-TRACKING instance. The orange vertices ensure that no cycle is untracked (see Figure 6).

■ **Figure 10** Illustration of a PLANAR-3-SAT instance (above) and the corresponding PLANAR-TRACKING instance associated with the reduction described in Section 4 (below).