



HAL
open science

Local certification of forbidden subgraphs

Nicolas Bousquet, Linda Cook, Laurent Feuilloley, Théo Pierron, Sébastien
Zeitoun

► **To cite this version:**

Nicolas Bousquet, Linda Cook, Laurent Feuilloley, Théo Pierron, Sébastien Zeitoun. Local certification of forbidden subgraphs. 2024. hal-04467821

HAL Id: hal-04467821

<https://hal.science/hal-04467821v1>

Preprint submitted on 20 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Local certification of forbidden subgraphs

NICOLAS BOUSQUET, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, France

LINDA COOK, Discrete Mathematics Group, Institute for Basic Science (IBS), Republic of Korea

LAURENT FEUILLOLEY, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, France

THÉO PIERRON, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, France

SÉBASTIEN ZEITOUN, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, France

Abstract

Detecting specific structures in a network has been a very active theme of research in distributed computing for at least a decade. In this paper, we start the study of subgraph detection from the perspective of local certification. Remember that a local certification is a distributed mechanism enabling the nodes of a network to check the correctness of the current configuration, thanks to small pieces of information called certificates. Our main question is: For a given graph H , what is the minimum certificate size that allows checking that the network does not contain H as a (possibly induced) subgraph?

We show a variety of lower and upper bounds, uncovering an interesting interplay between the optimal certificate size, the size of the forbidden subgraph, and the locality of the verification. Along the way we introduce several new technical tools, in particular what we call the *layered map*, which is not specific to forbidden subgraphs and that we expect to be useful for certifying many other properties.

1 INTRODUCTION

1.1 Context

Finding some given small structures in a graph, triangles for example, has become a major theme in the area of distributed graph algorithms. A lot of effort has been put recently in understanding various versions of this problem (detection, listing, counting, and testing) in several congested models, especially CONGEST and CONGESTED CLIQUE. We refer to the recent survey by Censor-Hillel [3] for an introduction to the topic and a full bibliography. This fundamental problem turned out to be very challenging and has been an incubator for new techniques, such as the expander decompositions [6] (that has later been used in other contexts, for example for derandomizing minimum spanning tree construction [7]).

Substructure detection has also become an important research topic in the more specific field of *local certification*. In local certification, one is interested in how much memory it takes to store a locally checkable certification of some given property. More precisely, for a given property, a local certification consists in a local algorithm taking as input a neighborhood around each node, along with labels assigned to nodes, called *certificates*, and outputting a binary decision, accept or reject. A certification scheme is correct if the following holds: there exists an assignment of certificates such that the local algorithm accepts at every node, if and only if, the property is satisfied in the network. Actually, the notion originates from self-stabilization [27], where one certifies the output of an algorithm, but in this paper we will focus on properties of the network itself. We refer to [16] for an introduction to local certification. (Formal definitions will be given later.)

Authors' addresses: Nicolas Bousquet, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, Villeurbanne, France, nicolas.bousquet@univ-lyon1.fr; Linda Cook, Discrete Mathematics Group, Institute for Basic Science (IBS), Daejeon, Republic of Korea, lindacook@ibs.re.kr; Laurent Feuilloley, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, Villeurbanne, France, laurent.feuilloy@univ-lyon1.fr; Théo Pierron, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, Villeurbanne, France, theo.pierron@univ-lyon1.fr; Sébastien Zeitoun, Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, Villeurbanne, France, sebastien.zeitoun@univ-lyon1.fr.

The classic measure of quality of a local certification is the size of the certificate, as a function of n , the size of the graph. For example, certifying that a graph is acyclic can be done with certificates of size $O(\log n)$, and this is optimal. One often separates graph properties into two broad categories: the ones that have a compact certification, that is of (poly)logarithmic size, and the ones that do not, and typically need linear or even quadratic size¹. This distinction can be seen as a local versus global dichotomy, if one considers the certificates to be the amount of global information needed to locally decide a property. Understanding what are the properties that fall in each of these regimes has been a major direction in the area in recent years, in particular with meta-theorems partially capturing the compact properties [17, 22, 23].

When it comes to local certification of network substructures, the positive case, that is certifying that a given substructure exists in the network, is often easy. Indeed, one can often simply put the names of the vertices involved in the certificates of all vertices, and use a locally-encoded spanning tree to point to a node of the structure. The real challenge is to certify that some structure is absent from the network, which is intuitively a much more global property. In that direction, the community has been focusing on one type of substructures, namely graph minors. Remember that a graph H is a minor of a graph G if it can be obtained from G by edge and vertex removals, as well as edge contractions. The following is one of the key questions of the area.

Conjecture 1.1 (Excluded-minor certification conjecture). *For every graph H , certifying that a graph is H -minor-free can be done with certificates of size $O_H(\log n)$.*

The conjecture has been proved for planar graphs [19], bounded-genus [13, 20], small excluded minors [1], path and cycle minors [17], and planar minors (with polylogarithmic certificates) [23]. A relaxed version has also been proved for every H [14], but the full conjecture is still widely open.

In this paper, our focus is not minors but another type of substructure, namely subgraphs, which are actually more popular in distributed computing, as discussed above. The general question we want to answer is the following.

Question 1.2. *What is the optimal label size for certifying that the network does not contain a fixed graph H as a non-induced or an induced subgraph?*

Remember that a graph H is a (non-induced) subgraph of G if it can be obtained from G by removing vertices and edges, and an induced subgraph of G if it can be obtained from G by removing only vertices (and the edges adjacent to these removed vertices).^{2,3} Subgraphs, which are very localized in the graph, might look easier to manage than minors, which can span a large part of the graph (because of the contraction operation). And indeed, if one allows the local verification algorithm to look at large enough (constant) distance, every node can see whether the forbidden subgraph appears or not in its neighborhood. But if we do not allow this, then the problem seems to become more challenging than for minors. Intuitively, since minors can appear in many ways in a graph, forbidding one constrains the graph structure a lot, and one can use this structure for certification. For example, forbidding a triangle as minor implies that the graph is a tree, while forbidding a triangle as induced subgraph still leaves a very complex graph class.

¹Any property can be certified with a quadratic number of bits in this model, with certificates encoding the adjacency matrix, see e.g. [16].

²In the literature, these two cases are simply called *subgraph* and *induced subgraph*, but to better differentiate we often add *non-induced* or *non-necessarily induced* to qualify the first case.

³For subgraph detection, induced subgraphs have been less popular, for no obvious reason, except maybe that they are difficult to manage. A paper that focuses on induced subgraphs is [28].

1.2 Discussion of previous work and refinement of the question

As mentioned earlier, there is a large literature on subgraphs detection in various distributed models. In general, results in these models are difficult to adapt to local certification. Indeed, on the one hand, the non-constructive aspect of certificates makes the latter a stronger model, but on the other hand, local certification is a broadcast-type model, in the sense that all neighbors see the same information from a specific node, which implies that comparisons with unicast models such as CONGEST are rarely useful.⁴

The only model that is clearly useful for us is BROADCAST CONGEST where, at each round, each node sends the same $O(\log n)$ -bit message to all its neighbors. Suppose that there is an upper bound of $f(n)$ rounds in this model to detect some subgraph (where detection means that if the graph contains one or several copies of the subgraph, then at least one node will detect at least one copy). Then we can derive an upper bound of $O(f(n) \log n)$ for certification, by encoding in the certificate of each node all the messages that it sends during the run of the algorithm. Indeed, given this information, every node can check that the run is correct, and a node rejects this certification in the end, if and only if, the output of the original algorithm is a detection.

A nice result proved in [15, 26] is that any tree on k vertices can be detected in $O(k2^k)$ rounds of the BROADCAST CONGEST model. From this result and the reasoning above, we derive the following theorem.

THEOREM 1.3 (DERIVED FROM [15, 26]). *For any tree T , certifying that a graph contains T as a (not necessarily induced) subgraph can be done with certificates of size $O(\log n)$ bits.*

Actually, one can also derive this result from the meta-theorem of [17] (see Corollary 2.7, noting that forbidding a path of length t as a minor is equivalent to forbidding it as a (non-induced) subgraph). From the same paper, one can adapt a matching lower bound (see Section 7 in [17]).

Given that the techniques to obtain Theorem 1.3 are very specific to non-induced subgraphs (roughly, a lot of information about non-edges is lost in the kind of compression that allows to store so little information), a first concrete question is the following.

Question 1.4. *What is the optimal certificate size for induced trees?*

For this question, the only non-trivial case known in the literature is for P_4 -free graphs⁵, which can be certified with $O(\log n)$ bits [22]. This proof is not really helpful to understand the general case, since it heavily uses the fact that P_4 -free graphs are the same as cographs, a graph class whose well-understood structure can be efficiently utilized for certification.

For lower bounds, results in BROADCAST CONGEST cannot be adapted directly, but the techniques can sometimes be borrowed. This is the case for the following theorem for cliques. (Note that, for cliques, induced and non-induced subgraphs are equivalent.)

THEOREM 1.5 (DERIVED FROM PROPOSITION 5 IN [9], SEE ALSO [11]). *For constant k , certifying K_k -free graphs requires certificates of size $\Omega(n/e^{O(\log \sqrt{n})})$ bits.*

Actually, [9] only mentions the result for triangles, but it is easy to generalize it to cliques. Indeed, if there would exist a more compact certification of larger cliques, one could use it to certify triangles, by simulating $k - 3$ virtual nodes adjacent to every node of the graph, *i.e.* by giving their certificate in a K_k -freeness certification to every node of the original graph.

At that point, one might simply conclude that being K_k -free is a global property, since it needs polynomial size proof, and move on to the next problem. We claim that understanding where is the

⁴The topic of broadcast versus unicast in local certification has been explored in [30].

⁵Remember that P_k denotes the path on k vertices, and we will also use K_k for cliques on k vertices.

right complexity on the spectrum between $\Theta(n)$ and $\Theta(n^2)$ (or more precisely $\Theta(m)$, where m is the number of edges) is actually an important question. A useful analogy here is with the CONGEST model, where algorithm running in $O(n)$ rounds are considered efficient (and are often optimal), whereas larger complexities, especially $\Theta(n^2)$, means that the problem is hard. The same intuition is true for local certification: using $o(n^2)$ bits means that (in dense graphs) we can avoid transferring information about all the edges, and having (quasi-)linear certificates means that one can store only a small piece of information per node of the graph. We will see later that this approach is also fruitful technically, since it forces one to design nice new tools and data structures.

We phrase this point of view into a question.

Question 1.6. *For the (induced) subgraphs that require polynomial certificate size, what are the right polynomials? Are they (quasi)-linear in the number of vertices, or in the number of edges? Or in between?*

As a side remark, note that when forbidding non necessarily induced subgraphs or minors, and taking only n as a parameter, one can use the fact that the number of edges is bounded away from n^2 , but this is not true for induced subgraphs (for example a clique does not contain any induced copy of P_k for $k > 2$).

For the case of cliques, we can actually almost match the lower bound of Theorem 1.5.

THEOREM 1.7. *Certifying that a graph does not contain a given constant-size clique as a subgraph can be done with $O(n \log n)$ bits.*

To establish this theorem, we give as a certificate to each node the list of the identifiers of its neighbors. Of course, this is useless for the node itself, since it knows this information, but it allows its neighbors to have knowledge at distance 2, which is enough for detecting cliques.⁶ Note that this techniques is not restricted to cliques: it applies to any setting where the forbidden subgraph would be detected automatically if we could increment the radius by 1.

After paths and cliques, a natural direction is to target cycles and some specific small subgraphs (e.g. diamonds). Here, we know from personal communication [8] that some cases have been solved: non-induced cycles require $\Omega(n)$ bits, and the technique based on triangles [9, 11] allows proving subpolynomial lower bounds for various small graphs. We take another direction, that opens new perspectives and, as side results, we will also obtain answers to some of these questions.

As mentioned earlier, certifying that a subgraph is absent from a graph is trivial if the nodes can see far enough in the graph, and the results above illustrate that we have some understanding of the distance 1 case. Our main question is: what happens in between?

Question 1.8. *What is the optimal size for forbidden induced or non-induced subgraphs, as a function of the verification radius and of the subgraph size?*

It is very unclear what one should expect. Is it the case that, even with radius 2, the certificate size drops? Or on the contrary, one should see (almost) all the subgraphs in order to gain from a larger verification view? Note that increasing the verification radius beyond 1 means that the view of a vertex might contain a quadratic number of edges. This rules out the intuitive lower bound argument that the view at distance 1 of a node contains only at most n edges and n certificates, when n^2 bits of information is necessary. In this sense, the lower bounds are more difficult, and adding on top that we look for constructions that hold even for induced subgraphs, makes it especially challenging. Also, the increased radius allows for more fancy upper bound techniques, for example if some node would “naturally” receive a large certificate, we can cut it into pieces and distribute it to its neighbors to average the maximum certificate size.

⁶To our knowledge, this observation has never been used in local certification, but can be considered as folklore.

Note that the idea of increasing the verification radius is not new: it has been introduced for constant radius (larger than 1) in [24], and non-constant radius in [18, 29], with recent new developments motivated by the so-called trade-off conjecture [2]. In this paper, the focus is slightly different, since the radius has to be compared to the subgraph size. In particular, questions like “can we go down to constant-size certificates if we look far enough?” are trivial to answer.

Additional discussion and motivation will be given in Section 1.4, after we describe our results.

1.3 Our results and techniques

The results of this paper are upper and lower bounds for certifying that a given subgraph does not appear in the network. The lower bound applies to any verification radius k , while the upper bounds apply for $k \geq 2$.

Paths as benchmark. For concreteness, let us start with Table 1, that summarizes the different results we get for induced paths.

Forbidden induced subgraph	Certificate size	Reference
P_{2k-1}	0	Direct
P_{2k+1}	$O(n \log n)$	Discussion after Theorem 1.7
P_{3k-1}	$O(n \log^3 n)$	Theorem 1.12
$P_{\lceil \frac{14}{3}k \rceil - 1}$	$O(n^{3/2} \log^2 n)$	Theorem 1.13
P_{4k+3}	$\Omega(\frac{n}{k})$	Theorem 1.9

Table 1. Bounds for induced paths verification when vertices can see at distance $k \geq 2$.

We will see later in this section that several of our results hold for more general subgraphs (induced and non-necessarily induced), but paths are good as a benchmark, and all our techniques were first designed for paths and then generalized.

Let us make a few observations on this table before we review the theorems and techniques. A first observation is that all of these bounds are in the polynomial regime. Hence a general insight is that compact certifications do not exist for forbidden subgraphs, except for specific cases (e.g. P_4 -free graphs and non-necessarily induced trees) or for cases where we forbid several subgraphs (see the additional related work, Subsection 1.4). But, as argued earlier, there is an important difference between certifying with $\Theta(n)$ bits and $\Theta(n^2)$ bits, and in this regime, a lot can be done. A second observation is that our upper bounds are increasing faster when we increase the ratio between the path length and the radius, but without crashing into $\Theta(n^2)$. This suggests that there might be a trade-off, although our lower bound does not allow such a parametrization. Finally, since our lower bound applies for paths of length at least $4k + 3$, we do not rule out that below this length a very compact certification exist. We do not believe this to be true, and think that the right bound should be close to our upper bounds. Also note that we do have non-trivial polynomial upper and lower bounds for lengths between $4k + 3$ and $\lceil (14/3)k \rceil - 1$.

Lower bounds. Our lower bound is the following.

THEOREM 1.9. *When the vertices can see at distance k , at least $\Omega(\frac{n}{k})$ bits are needed to certify T -free graphs, where T is either P_{4k+3} or a tree of diameter at least $4k + 2$ without degree 2 vertices.*

While we designed our lower bound arguments to handle paths, we believe they can be extended to all trees of diameter at least $4k + 2$, but the proof becomes more technical. To avoid too much technicalities, we rather extend it only to the trees without degree 2 vertices.

The general technique to establish this lower bound is inspired by the now classic reduction from non-deterministic communication complexity. There are actually several challenges to overcome for this to work with induced structures, but this discussion is deferred to Section 3.

Note that when restricting to radius 1, this theorem implies that P_7 -free graphs are hard to certify, while we already know that P_4 -free graphs are easy. We leave the remaining path lengths open.

Overview of our core upper bound technique. All our upper bounds use the same tool, that we call a *layered map*, which can be seen as a refinement of the so-called *universal scheme* (that we remind below). Just like the universal scheme, our layered maps are not specific to one type of task, and we believe they can be used in various contexts beyond subgraph finding. On the contrary, the tools we implement on top of these layered maps, are specific to subgraph detection.

Remember that in the universal scheme, the prover gives as certificate to every node the full map of the graph, and the vertices check that they have been given the same map as their neighbors, and that the map locally coincides with their neighborhoods. This takes $O(n \log n + \min(m \log n, n^2))$ bits (where m is the number of edges), using adjacency matrix or adjacency lists. Now, to make use of the larger verification radius, we use an idea from [18], which is to spread the map. More precisely: on correct instances, the prover cuts the map into pieces and distribute it to the nodes of the graph in such a way all the vertices can see all the pieces in their neighborhood. Every vertex then simply reconstructs the map, and resumes the universal scheme.

This technique works well if the neighborhood at distance k is large enough to drastically decrease the certificate size. This is the case for example when the minimum degree of the graph is polynomial, but in general this assumption is not met. Our layered map is a relaxed version of this spread universal scheme, where the nodes are given different certificates depending on their degree, and are able to check “partial maps” of the graph.

We describe the technique in the simpler setting with two groups of vertices, separating high and low degrees with a threshold of \sqrt{n} , and we show how this can be used to prove the following theorem (which does not appear in Table 1 since it is superseded by better results).

THEOREM 1.10. *We can certify P_{4k-1} -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

Intuitively, the high-degree vertices have enough neighbors to spread the whole map in their neighborhood, while the subgraph restricted to the low degree vertices is sparse enough to be given to all vertices (without spreading). The difficulty is that if the maps given to two high-degree vertices are not equal, and if these nodes are separated by enough low degree nodes, no node will be able to detect the inconsistency.

Actually, thanks to the radius k , if two high-degree vertices can be joined by hopping to high-degree vertices at distance at most $2k$, and the certification is accepted, then their maps must be identical. Therefore, we can define so-called *extended connected components*, or *ECC* for short, which partition the sets of high-degree vertices into groups that must have the same map. Our first technical work is to prove that we can certify exactly to all vertices the list of the ECCs, and how they partition the vertex set, with linear-size certificates.

We now sketch how to prove Theorem 1.10. Suppose that G contains a path P on $4k - 1$ vertices. Because different ECCs are at distance at least $2k$, such a path can touch zero, one or two ECCs. In the case of zero or one ECCs, the inconsistency issue raised earlier does not appear, and either the maps given are incorrect (and detected as incorrect), or they are correct and the vertices can see that they describe a graph with a long path. The case of two ECCs is more tricky. Essentially what we do is to give and certify for every node the length of the longest path starting from it and going towards the closest ECC. Along with other arguments, this ensures that at least one vertex will detect a long path touching two ECCs. (More details in Section 5.)

List of the forbidden subgraph upper bounds. It will be easier to sketch the proofs of our improved theorems once we will have formally described the layered maps, hence here we just list our results for reference, and refer to Section 5 for discussion of the proof techniques.

THEOREM 1.11. *For every $k \geq 2$, we can certify H -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$ for every H of size at most $4k - 1$.*

Actually, this theorem also holds for non-necessarily induced graphs. (The other theorems too, but it is not relevant since they are about paths, and we know Theorem 1.3).

THEOREM 1.12. *For every $k \geq 2$, we can certify P_{3k-1} -free graphs by looking at distance k with certificates of size $O(n \log^3 n)$.*

THEOREM 1.13. *For every $k \geq 2$, we can certify $P_{\lceil \frac{14}{3}k \rceil - 1}$ -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

1.4 Additional related work and discussions

Certification size landscape. A fruitful line of work in the LOCAL model consists in establishing the landscape of complexities for the classic family of problems called Locally checkable languages, LCLs for short; See e.g. [31]. In this perspective, one aims at characterizing the functions $f(n)$ for which there exists a problem whose optimal complexity is $f(n)$. Developing the same kind of theory for local certification size is an exciting research direction. So far, the only sizes for which we have natural problems with tight bounds are: $O(1)$, $\Theta(\log n)$, $\tilde{\Theta}(n)$ and $\tilde{\Theta}(m)$, where $\tilde{\Theta}$ means up to subpolynomial factors. (One can actually build a problem for any size $f(n, m)$ above $\log n$, but it is artificial in the sense that the definition of the problem refers to the function f .)

The constant certification size regime contains local properties such as coloring (intuitively the LCLs), and the dependency in other parameters such as the maximum degree has been explored very recently [2]. The size $\Theta(\log n)$ is very common, with the archetypal problem being acyclicity. See for example the list of problems in [24], and the recent series of papers establishing meta-theorems for this regime [17, 22, 23]. Actually, in some of these papers, the upper bounds are polylogs and not logs, and they are not matching the lower bounds, leaving open whether some of these properties have optimal certificate size $\Theta(\log^c n)$, with $c > 1$, or not.⁷ Finally, in the polynomial regime, we know problems with complexity $\tilde{\Theta}(n)$ (e.g. diameter ≤ 3 [5]) and problems with complexity $\tilde{\Theta}(m)$ (e.g. symmetric graphs [24] and non-3-colorable graphs [24]). Our paper explores the polynomial regime, but unfortunately since the bounds do not match we cannot conclude for sizes between n and m . Nevertheless, we believe that the case of induced paths of various sizes is a promising direction in this landscape perspective.

Certification of graph classes and forbidden subgraphs. A recent trend in local certification has been to focus on certifying graph property (e.g. planarity) instead of certifying the output of an algorithm (e.g. that a set of pointers distributed on the nodes collectively form a spanning tree). Two motivations behind this focus are that the optimal certification size can be seen as a measure the locality of a graph property (raising interest from the graph theory community) and that algorithms tailored to work on specific graph classes make more sense if one can ensure that the graph indeed belongs to the class.

Among the classes studied, many are defined by forbidden minors, as discussed after Conjecture 1.1. But some of these classes are not closed under minors, and are better described by families

⁷Minimum spanning trees with polynomial weights have optimal certificate size $\Theta(\log^2 n)$ [27], but one of the $\log n$'s originates from the encoding of the weights, hence this is not very satisfactory.

of forbidden induced subgraphs. For example, the authors of [25] tackle the case of chordal graphs (that are characterized by forbidding all induced cycles of length at least 4) and several other classes whose subgraph characterizations are more cumbersome: interval graphs, circular arc graphs, trapezoid graphs and permutation graphs. For all these classes, [25] establishes a $O(\log n)$ -bit certification. This might come as a surprise, since in this paper all the certifications are in the polynomial regime. The reason for this contrast is that all the classes we have just mentioned are very structured, and in particular have geometric representations, which can be used for certification. Very recently, it was proved that some classes are hard to certify even though they do have geometric representation: 1-planar graphs, unit disk graphs and other related classes require (quasi)-linear-in- n certificates [10].

Techniques based on bucketing by degree. As said earlier, our layered maps are based on bucketing the vertices by degree. This is a classic step in distributed subgraph detection. A classic canvas for detection algorithm is to first process the high-degree nodes, either arguing that there are few of them (see e.g. [21]) or that many nodes are close to high degree nodes (see e.g. [4]), and then to process low-degree nodes, often using color coding (a large part of these algorithms is randomized). Another technique consists in computing an expander (or conductance) decomposition, to separate the graph into parts that are well-connected (where one can basically compute as if the cluster would be a clique) and parts of low degree (see e.g. [6, 12]).

Our understanding is that only general intuitions can be transferred from these algorithms to our setting, namely: inside a cluster one should use the fact that it is easy to move/distribute information, whereas in low-degree parts one should enjoy the fact that there are fewer edges, hence fewer information to spread.

2 MODEL AND DEFINITIONS

2.1 Graph theory notions

In this paper, the network is modeled by an undirected graph without loops or parallel edges. The number of nodes is denoted by n and the number of edges is denoted by m . A graph is d -regular if all its vertices have degree d . A graph is *regular* if it is d -regular for some d .

We call H an *induced subgraph* of G if H is obtained from G by deleting a subset of the vertices of G and the edges incident to them. We call a graph H -free if it does not contain H as an induced subgraph. We call H a *subgraph* (sometimes specified as *non-necessarily induced*) of G if H is obtained from G by deleting a subset of the vertices of G , the edges incident to them, and an arbitrary subset of edges of G . We let P_k, C_k denote the path and cycle on k vertices respectively. The length of a path or cycle is the number of its edges.

We say two disjoint sets of vertices X, Y are *complete* to each other if all possible edges between X and Y are present. If there are no edges between X and Y we say they are *anticomplete*. A set of edges M of a bipartite graph G is a perfect matching if all the vertices of the graph are adjacent to exactly one edge of M . A set of edges in a bipartite graph is an *antimatching*, if they form a complete bipartite graph without a perfect matching.

2.2 Local certification

In the networks we consider, the vertices are equipped with unique identifiers on $O(\log n)$ bits. Certificates are labels attached to the vertices. The view at distance d of a vertex v consists of: (1) the vertices at distance at most d from v , (2) the identifiers and certificates of these vertices, and (3) the edges between these vertices, except the ones between two vertices at distance exactly d .

Definition 2.1. We say that there exists a local certification at distance d of size s for a property P if there exists a local algorithm (called the verification algorithm) taking as input on every node v the view at distance d , and outputting *accept/reject* such that:

- For every graph that satisfies the property P , there exists a certificate assignment, with certificates of size at most $s(n)$, such that the verification algorithm accepts at every node.
- For every graph that does not satisfy the property P , for all certificate assignments, there exists at least one node where the verification algorithm rejects.

In order to facilitate the writing, we say that the certificates are given by a *prover*. We can specify how the certificates are assigned by the prover on correct instances (i.e. graphs satisfying the property P), but we cannot control what happens on incorrect instances.

As an example, let us describe a local certification at distance 1, with certificate size $O(\log n)$, for checking that the graph is acyclic. On a correct instance, the prover chooses a node to be the root, and gives as a certificate to every node its distance to the root. The verification algorithm checks that the distances are consistent (typically that one neighbor has been assigned a strictly smaller distance, and the others a strictly larger distance). Now if the graph has a cycle, for any certificate assignment, the vertex with the largest assigned distance in the cycle has at least two vertices with distance smaller than or equal to its distance, hence it rejects.

3 LOWER BOUND FOR PATHS

In this section, we prove the following theorem.

THEOREM 1.9. *When the vertices can see at distance k , at least $\Omega(\frac{n}{k})$ bits are needed to certify T -free graphs, where T is either P_{4k+3} or a tree of diameter at least $4k + 2$ without degree 2 vertices.*

Discussion of the challenges and features of the proof. At first sight, the proof of this theorem is a reduction from the problem of non-disjointness in non-deterministic communication complexity, which is now a classic tool in the area. We would like to highlight several challenges and new features. First, we actually do not use communication complexity: we use a simple counting argument, avoiding defining all the non-deterministic Alice-Bob set-up, and demystifying the core of the argument. Second, the construction itself is very delicate. In classic proofs of the same flavor (e.g. [5, 18]), one usually deals with a radius k view in a very simple way: just replacing every edge with a path of length k . This is no good for us since we want to control very precisely the path structure of the graph. Also, we need to care about non-edges everywhere in the construction and not only in the part “encoding the disjointness”. Indeed, non-edges are crucial, since we know from Theorem 1.3 that for non-necessarily induced paths, such a lower bound cannot hold. These difficulties translate into a construction that blends cliques, matchings and antimatchings, and a proof that requires a careful case analysis to rule out unexpected paths.

PROOF. Let $k, n \in \mathbb{N}$, and T be either P_{4k+3} or a tree of diameter at least $4k + 2$ without degree 2 vertices. Let H, H' be two bipartite graphs on the same vertex set $\{0, 1\} \times \{1, \dots, n\}$. To have lighter notations, we denote an edge by (i, j) instead of $((0, i), (1, j))$. Assume moreover that H satisfies the two following properties:

- (i) For all $(i, j) \in \{1, \dots, n\}^2$, $(i, j) \in H$ if and only if $(j, i) \in H$ (H is symmetric)
- (ii) For all $i \in \{1, \dots, n\}$, $(i, i) \in H$ (H is reflexive)

and the same holds for H' . We will construct a graph $G_{k,n}(H, H')$ on $4k(n - 1) + |T|$ vertices, which is T -free if and only if H and H' do not have a common non-edge. The construction is illustrated on Figure 1.

Let us construct the graph $G_{k,n}(H, H')$ in the following way. We construct $4k$ cliques of size n , denoted by $K_1^0, \dots, K_{2k}^0, K_1^1, \dots, K_{2k}^1$. In each clique, we number the vertices from 1 to n .⁸ For every clique K and $i \in \{1, \dots, n\}$, we denote by $K[i]$ the vertex numbered by i in K . We add the following edges in $G_{k,n}(H, H')$:

- We put the bipartite graph H between K_1^0 and K_1^1 , and H' between K_{2k}^0 and K_{2k}^1 .⁹
- For every $j \in \{2, \dots, 2k-1\}$, we put the perfect matching $\{(K_j^0[i], K_j^1[i])\}_{1 \leq i \leq n}$.
- For every $j \in \{1, \dots, 2k-1\}$ and for every pair of cliques K, K' with $K \in \{K_j^0, K_j^1\}$, $K' \in \{K_{j+1}^0, K_{j+1}^1\}$, we put the antimatching defined by the complement of $\{(K[i], K'[i])\}_{1 \leq i \leq n}$ between K and K' (that is, all the edges but the matching $\{(K[i], K'[i])\}_{1 \leq i \leq n}$).

Finally, if $T = P_{4k+3}$, we add 3 more vertices v_0, v_1, w in the following way. The vertex v_0 (resp. v_1) is complete to K_1^0 (resp. K_1^1). The vertex w is complete to K_{2k}^0 and K_{2k}^1 . Otherwise, we choose a path P of length $4k$ in T containing no leaf, label its vertices by $v_1^0, \dots, v_{2k}^0, w, v_{2k}^1, \dots, v_1^1$, and for each connected component C of $T \setminus P$ adjacent with vertex v_a^b , we create a copy of C in $G_{k,n}(H, H')$ where the vertex adjacent with v_a^b becomes complete to K_a^b . Vertices in this copy are called *pending vertices* of K_a^b . We also add a copy of the connected component of w in $T \setminus E(P)$, and add all edges between w and $K_{2k}^0 \cup K_{2k}^1$.

Note that applying the latter construction to P_{4k+3} yields the former graph, hence we may use $G_{k,n}(H, H')$ in both cases.




Fig. 1. The graph $G_{k,n}(H, H')$. Each blob is a clique on n vertices. The double thick edges represent the graphs H and H' . The other double edges represent the matchings. The double dashed edges represent the antimatchings.

We state the main property satisfied by this construction in Proposition 3.1 below. We prove Proposition 3.1 by looking carefully at the structure of $G_{k,n}(H, H')$ at the end of this section.

Loosely speaking, the reason we require H and H' to be symmetric is because we want the cliques K_ℓ^0 and K_ℓ^1 to play the same “roles” in the graph for each $\ell \in \{1, 2, \dots, 2k\}$.

Proposition 3.1. *The graph $G_{k,n}(H, H')$ is T -free if and only if H and H' do not have a common non-edge (i, j) with $i, j \in \{1, 2, \dots, n\}$.²*

⁸We leave the precise unique identifier assignment implicit, for simplicity, but it is easy to create one, having the vertices of the first clique named 1 to n , then the second $n+1$ to $2n+1$, etc.

⁹That is, we identify the vertices of each side of the bipartite graph with the vertices of each blob, e.g. in increasing order of identifiers.

Let us prove Theorem 1.9 using Proposition 3.1. Assume that m bits are sufficient to certify T -free graphs when vertices can see at distance k . For any H, H' , let V denote the vertex set of $G_{k,n}(H, H')$.

We first define a family of positive instances. Let \overline{H} denote the bipartite graph obtained from H by replacing all edges (i, j) of H with non-edges and vice versa for any two *distinct* $i, j \in \{1, 2, \dots, n\}$. That is $E(\overline{H}) = (\{1, 2, \dots, n\}^2 \setminus E(H)) \cup \{(i, i)\}_{1 \leq i \leq n}$. Note that \overline{H} is symmetric and reflexive, and that H and \overline{H} have no common non-edge, by construction. Then, by Proposition 3.1, for every bipartite graph H which is symmetric and reflexive, the graph $G_{k,n}(H, \overline{H})$ is T -free. Therefore, by hypothesis, there exists a certificate function $c_H : V \rightarrow \{0, \dots, 2^m - 1\}$ such that all the vertices in $G_{k,n}(H, \overline{H})$ accept.

Claim 3.2. *For any two distinct symmetric and reflexive graphs H, H' on V , $c_H \neq c_{H'}$*

PROOF. Suppose that there are two different bipartite graphs H and H' satisfying the conditions above, such that $c_H = c_{H'}$. Since $H \neq H'$, we have either $E(H) \not\subseteq E(H')$ or $E(\overline{H}) \not\subseteq E(\overline{H'})$ or both. By symmetry, assume that $E(H) \not\subseteq E(H')$. Then, the graph $G_{k,n}(H', \overline{H})$ is accepted with c_H . Indeed, the vertices in the cliques K_j^i for $i \in \{0, 1\}$ and $j \in \{1, \dots, k\}$, or pending from these cliques have the same view as their view in $G_{k,n}(H', \overline{H'})$ since the only part of the graph which has changed between $G_{k,n}(H', \overline{H'})$ and $G_{k,n}(H', \overline{H})$ is the bipartite graph between K_{2k}^0 and K_{2k}^1 , which is at distance at least k from them. Thus, these vertices accept in $G_{k,n}(H', \overline{H})$. Similarly, w and the vertices in the cliques K_j^i for $i \in \{0, 1\}$ and $j \in \{k+1, \dots, 2k\}$ or pending from them have the same view in $G_{k,n}(H, \overline{H'})$ and $G_{k,n}(H', \overline{H})$, so they accept in the graph $G_{k,n}(H', \overline{H})$ as well. However, since $H \not\subseteq H'$, the graphs H' and $\overline{H'}$ have a common non-edge, hence $G_{k,n}(H', \overline{H'})$ is not T -free (by Proposition 3.1), which is a contradiction. ■

There are $2^{\frac{n(n-1)}{2}}$ such symmetric and reflexive bipartite graphs H on V , and there are at most $2^{m(4kn+|T|)}$ functions $V \rightarrow \{0, \dots, 2^m - 1\}$. Thus, by Claim 3.2, we get $2^{\frac{n(n-1)}{2}} \leq 2^{m(4kn+|T|)}$, which finally gives us $m = \Omega(\frac{n}{k})$. □

Let us now turn to the proof of Proposition 3.1. We will need the following result about the shape of long paths in $G_{k,n}(H, H')$. For succinctness, we will call $K_1^0, \dots, K_{2k}^0, K_1^1, \dots, K_{2k}^1$ the *cliques* (of $G_{k,n}(H, H')$) throughout this section.

Lemma 3.3. *Let T be an induced tree in $G_{k,n}(H, H')$. Then, T has at most $4k$ vertices in the cliques.*

PROOF OF LEMMA 3.3. The main ingredient of the proof of Lemma 3.3 is the following Claim 3.4. In the following, two cliques are said to be *antimatched* if there is an antimatching between them (that is, between the first and the second clique, there is a complete bipartite graph except for n independent non-edges).

Claim 3.4. *T has at most two vertices in each clique and*

- (i) *If T has two vertices in a clique K , then for all the cliques K' antimatched with K , T has at most one vertex in K' , and*
- (ii) *If T has two vertices in a clique K_j^i , then T has at most one vertex in K_j^{1-i} .*

PROOF. Since T is triangle-free, it contains at most two vertices in each clique. Suppose there is a clique $K := K_j^i$ and two distinct integers $x, y \in \{1, 2, \dots, n\}$ such that the vertices $K[x]$ and $K[y]$ are both in T . We prove (i) and (ii) separately:

- For (i), let K' be a clique antimatched to K . For each $z \in \{1, 2, \dots, n\}$, $K'[z]$ is adjacent to $K[x]$ (if $z \neq x$) or to $K[y]$ (if $z \neq y$). Thus, if there were at least two vertices in $T \cap K'$, T would have either a C_3 or a C_4 as a subgraph, which is a contradiction with T being a tree.

- We now prove property (ii). Let $K = K_j^i$ and $K' = K_j^{1-i}$. Assume for a contradiction that there exist two distinct $s, t \in \{1, 2, \dots, n\}$ such that $K'[s]$ and $K'[t]$ are both in T . Note that since T is triangle-free, if $j = 1$ then T cannot contain v_0 nor v_1 , and if $j = 2k$, then T cannot contain w . Therefore, since T has at least 5 vertices, it follows that there is some $K'' \in \{K_{j-1}^0, K_{j-1}^1, K_{j+1}^0, K_{j+1}^1\}$ such that T contains $K''[z]$ for some $z \in \{1, 2, \dots, n\}$. We have $x = z$ or $y = z$ otherwise, $K[x], K[y], K''[z]$ would be a triangle in T . Similarly, we have $s = z$ or $t = z$. Without loss of generality, we may assume that $x = s = z$. But then $K[z]-K[y]-K''[z]-K'[t]-K'[z]-K[z]$ is a C_5 subgraph with vertex set in $V(P)$, see Figure 2 for an illustration. (Note that the edge $K[z]-K'[z]$ arises even when $j \in \{1, 2k\}$ since H and H' are reflexive.) This is a contradiction with the fact that T is an induced tree. \square




Fig. 2. The case described in (ii) of Claim 3.4.

We are now ready to conclude the proof of Lemma 3.3. Assume by contradiction that T has at least $4k + 1$ vertices in the cliques. By the pigeonhole principle, there exists $j \in \{1, \dots, k\}$ such that T has at least 5 vertices in the cliques $K := K_{2j-1}^0, K' := K_{2j-1}^1, K'' := K_{2j}^0, K''' := K_{2j}^1$. Again by the pigeonhole principle, T has at least 2 vertices in one of these four cliques, and by Claim 3.4, exactly 2. By symmetry, assume that T has two vertices in K (see Figure 3 for an illustration). Then, by (i) and (ii) of Claim 3.4, T has at most one vertex in each of the cliques K', K'', K''' , so exactly one in each (because T contains five vertices in the union of these four cliques). Let us denote by $K[x], K[y]$ these two vertices in $K \cap T$, and by $K'[x'], K''[x''], K'''[x''']$ the vertices of T in K', K'', K''' respectively. Since T is an induced tree, we have $x'' = x$ or $x'' = y$ (else, $\{x, y, x''\}$ would be a triangle, since K'' is antimatched with K). By symmetry, we can assume that $x'' = x$. Similarly, we have $x''' = x$ or $x''' = y$. In the former case, $K[y]-K''[x]-K'''[x]-K[y]$ is a triangle in T . Hence, we have $x''' = y$.

Finally, there are two cases, depicted in Figure 3, which are the following:

- If $x' \in \{x, y\}$, assume by symmetry that $x' = x$. Then, $K[x]-K'[x]-K'''[y]-K[x]$ is a triangle in T , a contradiction.
- If $x' \notin \{x, y\}$, then T has a C_5 as a subgraph. This is a contradiction. \square

We are finally able to prove Proposition 3.1. We start by the case of paths, and then extend the proof to handle the other case.

PROOF OF PROPOSITION 3.1 FOR P_{4k+3} . Assume first that H and H' have a common non-edge. Let (i, j) be this common non-edge. By reflexivity of H and H' , we have $i \neq j$. Then, we can construct




Fig. 3. The two cases described at the end of the proof of Lemma 3.3.




Fig. 4. The path on $4k + 3$ vertices in $G_{k,n}(H, H')$ if H and H' have a common non-edge (i, j) .

an induced path of length $4k + 3$ as follows: $v_0, K_1^0[i], K_2^0[j], K_3^0[i], \dots, K_{2k}^0[j]$ (alternating between i and j), $w, K_{2k}^1[i], K_{2k-1}^1[j], K_{2k-2}^1[i], \dots, K_1^1[j], v_1$. This path is depicted on Figure 4. Note that since (i, j) is a non-edge in H' which is symmetric, then (j, i) is also a non-edge in H' . Thus, it is indeed an induced path.

Conversely, assume that there is an induced path P on $4k + 3$ vertices in $G_{n,k}(H, H')$, and let us show that H, H' have a common non-edge. By Lemma 3.3, P has at most $4k$ vertices in the cliques, and since there are exactly 3 vertices outside the cliques, P has exactly $4k$ vertices in the cliques, and it uses all the vertices outside of the cliques, namely v_0, v_1, w . Note that the neighborhoods of v_0 and v_1 in $G_{n,k}(H, H')$ are cliques, hence they must have degree 1 in P so they are the two endpoints of P . So P starts with $v_0, x_1 \in K_1^0$, and ends with $v_1, y_1 \in K_1^1$.

Since $w \in V(P)$ is not an end of P it must have degree two in P . Hence, for some $i, j \in \{1, 2, \dots, n\}$, $K_{2k}^0[j], K_{2k}^1[i]$ are the neighbors of w in $V(P)$. $K_{2k}^0[i], K_{2k}^1[j]$ must be non-adjacent, so (j, i) is not an edge of H' . Since H is reflexive it follows that $i \neq j$ and since it is symmetric (i, j) is not an edge of H' .

Claim 3.5. For each $\ell \in \{1, 2, \dots, 2k\}$ the vertices of P in $K_\ell^0 \cup K_\ell^1$ are exactly $K_\ell^0[i], K_\ell^1[j]$ or they are exactly $K_\ell^0[j], K_\ell^1[i]$.

PROOF. Let $\ell \in \{1, 2, \dots, 2k-1\}$ be maximum such that the claim does not hold for ℓ . By symmetry, we may assume that both $x := K_{\ell+1}^0[j], y := K_{\ell+1}^1[i]$ are elements of $V(P)$. Since v_0, v_1 are the ends of P , both x, y must have a neighbor in $K_\ell^0 \cup K_\ell^1$. In particular, there must be vertices $x', y' \in K_\ell^0 \cup K_\ell^1$ so that $(x, x'), (y, y')$ and there are no further edges between vertices in $\{x, x', y, y'\}$. By construction for each $K \in \{K_\ell^0, K_\ell^1\}$ and $K' \in \{K_{\ell+1}^0, K_{\ell+1}^1\}$, there is an antimatching between K and K' . So in particular, either $x' = K_\ell^0[i]$ and $y' = K_\ell^1[j]$ or $x' = K_\ell^1[i]$ and $y' = K_\ell^0[j]$. ■

Thus by Claim 3.5 and since H is symmetric, (i, j) is a non-edge of H . So H and H' have a common non-edge as desired. □

We may now adapt this proof to handle the case of trees without vertices of degree 2.

PROOF OF THE REMAINING CASE OF PROPOSITION 3.1. If H and H' share a common non-edge, then by the previous proof, there is an induced path on $4k + 3$ vertices in $G_{k,n}(H, H')$ as depicted on Figure 4. Adding all the vertices outside of the cliques to this path yields a copy of T .

Conversely, assume that there is an induced copy of T in $G_{n,k}(H, H')$, and let us show that H, H' have a common non-edge. By Lemma 3.3, T has at most $4k$ vertices in the cliques, and since there are exactly $|V(T)| - 4k$ vertices outside the cliques, T has exactly $4k$ vertices in the cliques, and it uses all the vertices outside of the cliques. In particular, by connectivity, T must contain at least a vertex in each clique (hence exactly once). Using this, we can show that Claim 3.5 still holds in this case, which concludes. □

4 UPPER BOUND CORE TECHNIQUE: LAYERED MAPS

4.1 Spread universal certification and graphs of large minimum degree

We have described in Subsection 1.3 the general idea of layered maps, which are a refinement of the technique of spreading the certificate of the universal scheme. We formalize the spread universal scheme (introduced in [18]) in the following theorem.

- THEOREM 4.1.** (i) *Let $\delta < 1$. Any property can be certified with certificates of size $O(n^{2-\delta} \log n)$ on graphs of minimum degree n^δ if vertices can see at distance 2.*
(ii) *Any property can be certified with certificates of size $O(n \log^2 n)$ on regular graphs if vertices can see at distance 2.*

Before giving the proof, we give a short sketch for the case of the first item. On correct instances, the prover computes the certificate of the universal scheme in the form of an adjacency matrix, cuts it into n^δ pieces of size $n^{2-\delta}$, and gives $O(\log n)$ pieces to every node, in such a way that each vertex can see all the pieces in its neighborhood. Such an assignment of pieces is proved to exist via probabilistic method and coupon collector. Note that we need the distance to be at least two, only to ensure that every node can check that the graph it reconstructs is the same as its neighbors.

Before proving Theorem 4.1, we show the following Lemma 4.2 (inspired from the coupon collector theorem), that we will use several times.

Lemma 4.2. *Let G be a n -vertex graph and $0 < d < n$. Let \mathcal{P} be a set of size d , called the set of pieces. Then, there exists a way to assign $3 \log n$ pieces to every vertex of G such that, for every vertex u of degree at least d , each piece of \mathcal{P} has been given to u or one of its neighbors.*

PROOF. Let us consider a random assignation of the pieces to the vertices of G , where each vertex gets assigned a subset of \mathcal{P} of size $3 \log n$ pieces (uniformly at random, independently of all the other vertices). Let us fix some piece $p \in \mathcal{P}$ and some vertex u of degree at least d . Observe that

the probability that p has not been given to u not to its neighbors is at most:

$$\left(1 - \frac{3 \log n}{d}\right)^d \leq e^{-3 \log n} \leq \frac{1}{n^3}.$$

By union bound, the probability that such an event occurs for some vertex u of degree at least d and some piece $p \in \mathcal{P}$ is at most $\frac{1}{n}$. This is strictly smaller than 1 for all $n \geq 2$. Thus, there exists some correct assignment of the pieces. \square

We may now proceed with the proof of Theorem 4.1.

PROOF OF THEOREM 4.1. (i) Let us describe a certification scheme of size $O(n^{2-\delta} \log n)$ for any property, in graphs of minimum degree n^δ , if vertices can see at distance 2. Let $G = (V, E)$ be a graph with $|V| = n$, and minimum degree n^δ .

Certification. The prover assigns to each vertex u a certificate divided in two fields: one common to every node (denoted by `SpanningTree`), and another specific to u (denoted by `Pieces(u)`).

- In `SpanningTree`, the prover writes a spanning tree \mathcal{T} of G , that is: the identifiers of the vertices and of their parent in \mathcal{T} . This uses $O(n \log n)$ bits.
- For `Pieces`, the prover cuts the adjacency matrix of G in n^δ parts, each of size $n^{2-\delta}$, and number these parts from 1 to n^δ . Then, using Lemma 4.2, since all the vertices have degree at least n^δ , it may write $3 \log n$ numbered pieces in each `Pieces(u)` such that every vertex sees each piece at least once in its closed neighborhood. This certificate has size $O(n^{2-\delta} \log n)$.

Verification. The verification algorithm of each vertex u is done in two steps.

- The first step is to check the correctness of `SpanningTree`. To do so, u checks if `SpanningTree` is the same in its certificate and the certificate of all its neighbors. Then, it checks if it is indeed a tree, and if each of its neighbors in `SpanningTree` is also a neighbor in G . If it is not the case, u rejects.

If no vertex rejects at this point, `SpanningTree` is a correct spanning tree known by all the vertices. In particular, every vertex u knows n , the number of vertices in G .

- The second and main step of the verification is the following. Since u knows n , u knows the number pieces (which is n^δ). The vertex u checks if it sees each numbered piece appears in `Pieces(v)` for some vertex v in its closed neighborhood, and rejects if it is not the case. If u did not reject, it may thus reconstruct the whole graph G , and can also recover the graph reconstructed by every of its neighbors since it sees at distance at least 2. If some of these reconstructed graphs are different, u rejects. Then, u checks if its neighborhood is correctly written in its reconstructed graph. Finally, if u did not reject before, it accepts if and only if its reconstructed graph satisfies the property.

Correctness. If no vertex rejects in the verification procedure, all the vertices reconstructed the same graph. Moreover, this reconstructed graph is equal to G , because every vertex checked that its neighborhood is correct in it. Since every vertex accepts if and only if this graph satisfies the property, the scheme is correct.

- (ii) Let us now describe a certification scheme of size $O(n \log^2 n)$ for any property in regular graphs, if vertices can see at distance 2. Let G be an n -vertex d -regular graph.

Certification. The certification is similar to the proof of (i), except that since G is a d -regular graph, it can be represented with $O(dn \log n)$ bits using adjacency lists (instead of the adjacency matrix). The prover cuts it in d pieces, each of size $O(n \log n)$, numbers them

from 1 to d , and applies Lemma 4.2 to assign $3 \log n$ pieces in the certificate of every vertex. This yields certificates of size $O(n \log^2 n)$.

Verification. The verification algorithm is the following one. Every vertex u knows its own degree, which is equal to the number of pieces. So u can check if it sees all the pieces in its closed neighborhood. Then, as for (i), u reconstructs the whole graph G and checks if all its neighbors reconstructed the same graph. Finally, u checks if its own neighborhood is correctly written, and if the property is satisfied.

Correctness. As for (i), if no vertex rejects in the verification procedure, all the vertices computed the same graph, which is equal to G and satisfies the property. \square

4.2 Definitions: extended connected components, ECC table, and more

We now introduce formally the layered map, that we have informally described in the introduction, along with other related notions.

4.2.1 Definitions.

Extended connected components. Let $\varepsilon > 0$ and $N = \lceil \frac{1}{\varepsilon} \rceil$. We consider the following partition of the vertices depending on their degrees: for every $i \in \{1, \dots, N\}$, let $V_i := \{u \in V \mid n^{(i-1)\varepsilon} \leq \deg(u) < n^{i\varepsilon}\}$ (where $\deg(u)$ is the degree of the vertex u). Note that the sets V_1, \dots, V_N form a partition of V . Let $L_i := \bigcup_{j \leq i} V_j$, that is, L_i is the set of vertices having degree less than $n^{i\varepsilon}$. And let $H_i := \bigcup_{j \geq i} V_j$, that is, H_i is the set of vertices having degree at least $n^{(i-1)\varepsilon}$ (it is also $V \setminus L_{i-1}$). Note that $H_1 = L_N = V$.

Let $k \geq 2$, the distance up to which the vertices can see in the certification scheme. For each $i \in \{1, \dots, N\}$, we partition H_i in subsets called *extended connected components of H_i* (abbreviated by ECC_i). For $u, v \in H_i$ we say u, v are *i -linked* if there is a u, v path in G which does not have $2k - 2$ consecutive vertices in L_{i-1} (in other words, such a path should regularly contain vertices in H_i). An ECC_i is a maximal subset of vertices of H_i which are pairwise i -linked. Note that being i -linked is an equivalence relation. By definition, two vertices in two different ECC_i 's are at distance at least $2k - 1$ from each other in G .

Let us denote by \mathcal{E}_i the set of all ECC_i 's (which is a partition of H_i). For each $C \in \mathcal{E}_i$, we define the *identifier of C* , denoted by $\text{id}(C)$, as the smallest identifier of a vertex in C . Finally, we denote by $C(d)$ the set of vertices at distance exactly d from C (by convention, we set $C(0) := C$). Note that $C(0) \subseteq H_i$, and for all $d \geq 1$, we have $C(d) \subseteq L_{i-1} = V(G) \setminus H_i$.

Note that for every $u \notin H_i$ there is at most one $\text{ECC}_i C \in \mathcal{E}_i$ such that u is at distance at most $k - 1$ from C .

ECC-table. Let T_G be the following table, with n rows and N columns, called the *ECC-table of G* . The rows are indexed by the identifiers of the vertices, and the columns by $\{1, \dots, N\}$. Let $u \in V$ and $i \in \{1, \dots, N\}$. Let us describe the entry $T_G[u, i]$.

- If u is at distance at most $k - 1$ from some $\text{ECC}_i C_u$, we set $T_G[u, i]$ as $(\text{id}(C_u), d_u)$ where d_u is the distance from u to C_u . Note that $u \in H_i$ if and only if $d_u = 0$.
- Otherwise, we set $T_G[u, i] = \perp$.

In other words, for each $i \in \{1, \dots, N\}$, the partition of H_i in ECC_i 's is written in the table T_G , and for every vertex $u \in L_{i-1}$ "close" to some $C_u \in \mathcal{E}_i$, the distance from u to C_u is also stored in T_G .

Witnessed graph. We define a relation \leq on V in the following way. Let $u, v \in V$. We say that $v \leq u$ if there exists $i \in \{1, \dots, N\}$ such that $v \in V_i$, $u \in H_i$, and u, v are in the same ECC_i . Note that, by definition, \leq is transitive. For every $u \in V$, let $V_{\leq u}$ be the set $\{v \in V \mid v \leq u\}$, called the *set of vertices witnessed by u* . Let $G_{\leq u}$ be the subgraph of G obtained by keeping only edges having at

least one endpoint in $V_{\leq u}$, which is called the *graph witnessed by u* . Intuitively, if $v \in V_{\leq u}$, then u can check all the adjacencies of v thanks to the certificates (see Theorem 4.5).

Local computation scheme. Now, let us introduce the notion of *local computation scheme*. Informally, it is a tool allowing the vertices to perform a pre-computation at the beginning of their verification in a certification scheme¹⁰. We will use it so that each vertex u will pre-compute the ECC-table T_G and its witnessed graph $G_{\leq u}$.

Definition 4.3 (Local computation scheme). *Let f be a function taking as input a graph G and a vertex u of G . A local computation scheme for f of size s is a scheme where the prover gives certificates of size s to the vertices, and each vertex either rejects, or outputs something. We also require the two following soundness conditions to be satisfied for every graph $G = (V, E)$:*

- (i) *if no vertex rejects, then every vertex $u \in V$ outputs $f(G, u)$;*
- (ii) *there exists a certificate assignment such that no vertex rejects.*

4.3 Technical results

We now design local computation schemes to compute the ECC-table and the witnessed graphs, as stated in the following Theorems 4.4 and 4.5.

THEOREM 4.4. *There exists $c > 0$ such that, for all $0 < \varepsilon < 1$, there exists a local computation scheme for $f(G, u) := T_G$ of size $\frac{c}{\varepsilon} \cdot n \log n$, if the vertices can see at distance k .*

Note that the output of the function f in Theorem 4.4 does not depend on u , but only on G and ε .

PROOF. Let $0 < \varepsilon < 1$ and $N = \lceil \frac{1}{\varepsilon} \rceil$. Let us describe a local computation scheme for $f(G, u) = T_G$ of size $O(\frac{n}{\varepsilon} \log n)$.

Certification. The certificates of the vertices consists of information stored in three fields, denoted by `SpanningTree`, `Table` and `Components`. The certificate is the same for all the vertices.

- The prover chooses a spanning tree \mathcal{T} of G and writes it in `SpanningTree`. More precisely: it writes the identifiers of all the vertices and of their parents in \mathcal{T} using $O(n \log n)$ bits.
- In `Table`, the prover writes the table T_G . Since it has n rows, N columns, and $O(\log n)$ bits per cell, it has size $O(Nn \log n)$.
- In `Components`, the prover gives information to the nodes to check the correctness of the partition in ECC_i 's written in `Table`. For each $i \in \{1, \dots, N\}$, and for each $C \in \mathcal{E}_i$, the prover constructs the graph G_C , where the vertices of G_C are the vertices of C , and there is an edge between two vertices in G_C if and only if they are at distance at most $2k - 2$ in G . By definition of an ECC_i , G_C is connected. Then, the prover chooses a spanning tree \mathcal{T}_C of G_C and writes its structure in `Components`, with the identifier of the corresponding vertices. For each edge (u, v) in \mathcal{T}_C , there exists a vertex $w \in V$ at distance at most $k - 1$ from both u and v in G . The prover labels the edge (u, v) in \mathcal{T}_C by the identifier of w . For a given $C \in \mathcal{E}_i$, $O(|C| \log n)$ bits are required. In total, since the prover does this for each $C \in \mathcal{E}_i$ and every $i \in \{1, \dots, N\}$, it uses $O(Nn \log n)$ bits.

Since $N = \lceil \frac{1}{\varepsilon} \rceil$, the overall size of the certificate is thus $O(\frac{n}{\varepsilon} \log n)$.

Verification. The vertices perform the following verification procedure.

- (i) First, every vertex checks that its certificate is the same as the certificates of its neighbors, and rejects if it is not the case.

¹⁰This new definition allows us to compose different computations using certificates, which is not possible in general with the standard certification definition, because of its binary output.

- (ii) To check the correctness of `SpanningTree`, every vertex checks if it is indeed a tree, and if each of its neighbors written in `SpanningTree` is indeed a neighbor in G . It rejects if it is not the case.

If no vertex rejects at this point, the spanning tree of G written in `SpanningTree` is correct. In particular, all the vertices know n , and the whole set of the identifiers of vertices in G .

- (iii) The next two steps of the verification consist in checking if the partition in ECC_i 's written in T_G is correct. Every vertex u checks that for all $i \in \{1, \dots, N\}$, and for all vertices $v, w \in H_i$ at distance at most $k - 1$ from u , we have $\text{Table}[v, i] = \text{Table}[w, i]$, and if this common value is of the form $(id(C), 0)$. If it is not the case, u rejects.

If no vertex rejects at this point, then for all $i \in \{1, \dots, N\}$ and for every pair of vertices $v, w \in H_i$ which are at distance at most $2k - 2$, v and w are written to be in the same ECC_i in `Table`. By transitivity: if two vertices are in the same ECC_i in G , then they appear in the same ECC_i in `Table`.

- (iv) Every vertex $u \in V$ determines the index $j \in \{1, \dots, N\}$ such that $u \in V_j$. For every $i \leq j$, u does the following. We have $u \in H_i$. Let $C_u \in \mathcal{E}_i$ be the ECC_i of u . The vertex u checks that \mathcal{T}_{C_u} written in `Components` is indeed a tree. Moreover, for each edge (u, v) in \mathcal{T}_{C_u} labeled with the identifier of a vertex w , u checks if it sees indeed w at distance at most $k - 1$ (and rejects if it is not the case).

If no vertex rejects at this point, for all vertices $u, v \in V$, if $\text{Table}[u, i] = \text{Table}[v, i]$ and if this value is of the form $(id(C), 0)$, then u and v are indeed two vertices of H_i which are in the same ECC_i (indeed, \mathcal{T}_C is connected, so there exists a path from u to v in \mathcal{T}_C , which corresponds to a path in G which does not have $2k - 1$ consecutive vertices in L_{i-1}).

Thus, together with step (iii) of the verification, the partition in ECC_i 's written in `Table` is correct.

- (v) The next step of the verification consists in checking if the distances written in `Table` are also correct. To do so, every vertex u does the following. For each $i \in \{1, \dots, n\}$, if $u \notin H_i$, u checks that $\text{Table}[u, i] = \perp$ if and only if u does not see any vertex of H_i at distance at most $k - 1$. And if u sees vertices in some $C_u \in \mathcal{E}_i$ at distance d_u , u checks that $\text{Table}[u, i] = (id(C_u), d_u)$. If it is not the case, u rejects.

Computation. If a vertex u did not reject during the verification phase, it outputs `Table`. If no vertex rejects, we have $\text{Table} = T_G$ (since both the partition in ECC_i 's and the distances to the ECC_i 's written in `Table` are correct). Moreover, if the prover gives the certificates as described above, no vertex will reject. Thus, the computation scheme is correct. \square

THEOREM 4.5. *There exists $c > 0$ such that, for all $0 < \varepsilon < 1$, there exists a local computation scheme for $f(G, u) := G_{\leq u}$ of size $\frac{c}{\varepsilon} \cdot n^{1+\varepsilon} \cdot \log^2 n$, if the vertices can see at distance 2.*

PROOF OF THEOREM 4.5. Let $0 < \varepsilon < 1$ and $N = \lceil \frac{1}{\varepsilon} \rceil$. We describe the local computation scheme.¹¹

Certification. Let us describe the certificates given by the prover to the vertices on correct instances. First, it gives to every vertex u its certificate in the local computation scheme for T_G given by Theorem 4.4, which has size $O(\frac{n}{\varepsilon} \log n)$. Then, the prover gives other additional information to u , denoted by $\text{Pieces}(u)$, which is the following. For every $i \in \{1, \dots, N\}$, the prover constructs the graph G_i obtained from G by keeping only edges having at least one endpoint in L_i . Since each vertex in L_i has degree at most $n^{i\varepsilon}$, G_i has at most $n^{1+i\varepsilon}$ edges. Then, the prover cuts the adjacency

¹¹Note that in theorem the verification radius is 2, but this result is just a component of a scheme with radius $k \geq 2$, hence the ECC are still defined with respect to this general k .

list of G_i (which has size $n^{1+i\epsilon} \log n$) in $n^{(i-1)\epsilon}$ pieces of size $n^{1+\epsilon} \log n$, and number them from 1 to $n^{(i-1)\epsilon}$. Finally, for every vertex u , it writes $3 \log n$ numbered pieces of G_i in $\text{Pieces}(u)$ in such a way: for each $u \in H_i$, u sees every piece P of G_i in the certificate of at least one of its neighbors. Since each vertex $u \in H_i$ has degree at least $n^{(i-1)\epsilon}$, and since there are $n^{(i-1)\epsilon}$ different pieces of G_i , this is possible (see Lemma 4.2). In total, since the prover does this for every $i \in \{1, \dots, N\}$, $\text{Pieces}(u)$ has size $O(Nn^{1+\epsilon} \log^2 n)$, so the size of the overall certificate is $O(\frac{n^{1+\epsilon}}{\epsilon} \log^2 n)$.

Verification. The vertices perform the following verification procedure.

- (i) First, each vertex applies the verification of the local computation scheme for T_G given by Theorem 4.4.
- (ii) To verify $\text{Pieces}(u)$, each vertex u does the following: if it sees two numbered pieces $P \in \text{Pieces}(v)$ and $P' \in \text{Pieces}(w)$ for some v, w at distance at most $k - 1$, such that P and P' are two pieces of G_j for some $j \in \{1, \dots, N\}$ which are numbered the same, u checks that these two pieces are indeed the same. If it is not the case, u rejects.
- (iii) Finally, for every $i \in \{1, \dots, N\}$, each $u \in V_i$ does the following verification. For every $j \leq i$, and for every numbered piece P of G_j , u checks if it sees P in a certificate $\text{Pieces}(v)$ for some v in its closed neighborhood. If it is not the case, u rejects. And if u sees that its own edges and non-edges are not correctly written in the piece P of G_i where it should be, it rejects.

Computation. Let us describe the computation of every vertex u . If no vertex rejected at step (i), u computed the ECC-table T_G . In particular, u knows the partition of H_i in ECC_i 's for every $i \in \{1, \dots, N\}$. So u can compute the set $V_{\leq u}$.

Claim 4.6. *Let $u, v \in V$ such that $v \leq u$. If no vertex rejected during the verification phase, then u can compute all the edges and non-edges having v as an endpoint.*

PROOF. Let $i \in \{1, \dots, N\}$ such that $v \in V_i$ and $u \in H_i$. Because of step (iii) of the verification, the vertex u sees all the parts of G_i in its certificate or in the certificate of some of its neighbors. However, nothing ensures that these parts are correct. But thanks to step (ii), any two vertices in H_i at distance at most $2k - 2$ get the same pieces of G_i in their neighborhoods. So by transitivity, two vertices which are in the same ECC_i have the same pieces of G_i in their neighborhood. In particular, since u and v are in the same ECC_i , and since the edges and non-edges of v are correctly written in the pieces which are in the closed neighborhood of v (otherwise it would have rejected at step (iii)), then u knows the edges and non-edges of v . \square

Thus, if no vertex rejects in the verification phase, thanks to Claim 4.6, every vertex u is able to reconstruct the graph $G_{\leq u}$ thanks to the certificates Pieces it sees its neighborhood. Moreover, if the prover gives the certificates as described above, no vertex will reject. So the computation scheme is correct. \square

Remark 4.7. Note that, with the certificates given in the certification scheme of Theorem 4.5, each vertex v can also output the witnessed graph $G_{\leq u}$ of any vertex u at distance at most $k - 1$ from itself. Indeed, v can see all the pieces of $G_{\leq u}$ which are spread in certificates of vertices in $N[u]$.

5 OVERVIEW OF THE FORBIDDEN SUBGRAPH CERTIFICATIONS

In this section, we sketch how we use the layered maps for forbidden subgraphs certification. All the formal proofs are in Section 6.

We start with Theorem 1.10, that we have already discussed in the introduction, but now with the right notations, and more details. The statement is actually generalized in several ways later on, but it is good to discuss it first, to introduce the techniques one by one.

THEOREM 1.10. *We can certify P_{4k-1} -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

What the prover does is the following: first, it uses the computation schemes of Theorems 4.4 and 4.5, with $\varepsilon = \frac{1}{2}$, that is $N = 2$. Thus, V is divided into V_1 (the set of vertices having degree less than \sqrt{n} , the *low-degree vertices*) and V_2 (the set of vertices having degree higher than \sqrt{n} , the *high-degree vertices*). Note that $H_1 = V$ so there is only one ECC_1 which is the whole set V , and $H_2 = V_2$. With the local certification schemes of Theorems 4.4 and 4.5, every vertex u will compute its witnessed graph $G_{\leq u}$ and the ECC-table T_G . Let us recall that T_G gives the partition of V_2 in ECC_2 's, and also the distance from every vertex of V_1 to its closest ECC_2 (if it is at most $k - 1$).

For every vertex $v \in V_1$ and $u \in V$, we have $v \leq u$, hence u computes¹² all the edges and non-edges of v . In other words: the adjacency of low-degree vertices is known by every vertex. In particular, if G has an induced P_{4k-1} which is made only of small-degree vertices in G , it is detected by every vertex. Note also that for all $u, v \in V_2$ which are in the same ECC_2 , we have $v \leq u$, so u knows all the edges and non-edges of v . In particular, if there is an induced P_{4k-1} in G containing high-degree vertices from only one ECC_2 , any vertex u in this ECC_2 will detect it. So the difficult case is to detect an induced P_{4k-1} in G containing high-degree vertices from at least two different ECC_2 's. Using the fact that two distinct ECC_2 's are at distance at least $2k - 1$ from each other, we will see that it is in fact sufficient to consider the case of paths having big-degree vertices in exactly two different ECC_2 's.

To enable the vertices to detect the existence of paths of length $4k - 1$ going through exactly two distinct ECC_2 's, the prover will add some information in the certificates, which is the following. For each $v \in V$, if v is close to an ECC_2 denoted by C_v , let us denote by P_v the longest induced path which starts at v and has all its other vertices strictly closer from C_v than v . For every $v \in V$, the prover will give the length of P_v to all the vertices. Then, to verify the certificates, for every $C \in \mathcal{E}_2$, every vertex $u \in C$ will do the following: u checks if there exists a path in $G_{\leq u}$ which starts from u , ends in some vertex v close to some ECC_2 $C_v \neq C$, and can be extended into a path of length $4k - 1$ using the path P_v . We prove that if a P_{4k-1} path exists in the graph, touching two ECC_2 , then such a pair u, v exists.

A first extension of this theorem is to certify *any* forbidden subgraph on $4k - 1$ vertices.

THEOREM 1.11. *For every $k \geq 2$, we can certify H -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$ for every H of size at most $4k - 1$.*

To prove this, we will have to certify more information than the lengths of longest paths P_v we had before. Informally speaking, for every vertex v , we will need to determine, for every possible node h in H , if v can be node h of a copy of H that is in between two ECC_2 's.

Then we push the technique to get quasi-linear size certificates.

THEOREM 1.12. *For every $k \geq 2$, we can certify P_{3k-1} -free graphs by looking at distance k with certificates of size $O(n \log^3 n)$.*

To prove this, we will use the whole power of the witnessed graphs and layered maps since we will use it with an arbitrary number of layers (while the previous bounds only used two layers). Doing it introduces technicalities that did not appear in the previous cases (and that explain why we were only able to get P_{3k} -freeness instead of P_{4k-1} -freeness).

Finally, we achieve to increase the length of the path, keeping the same size as in Theorem 1.10.

¹²In the sense of Definition 4.3.

THEOREM 1.13. *For every $k \geq 2$, we can certify $P_{\lceil \frac{14}{3}k \rceil - 1}$ -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

Again the proof is inspired by the proof of Theorem 1.10. However, in order to get the result, we need to add different information in the certificate of each vertex. Indeed, so far, the certificate where almost global in the sense that, except in order to certify the graph $G_{\leq u}$ for every u in Theorem 4.5, the certificate given to vertices where the same. In that proof, we need to add local certificate on each node and then we will have to be more careful on which vertex can detect the existence of a path if a path exists. Very informally speaking, the information we will be able to store is, for every vertex v , what happens if only two vertices of a path belong to $C(d)$ (assuming that v is in $C(d)$).

6 FORMAL PROOFS OF THE FORBIDDEN SUBGRAPH CERTIFICATIONS

In this section, we prove all our upper bounds: Theorem 1.10, Theorem 1.11, Theorem 1.12 and Theorem 1.13. These are presented by increasing level of technicality.

6.1 Upper bound for P_{4k-1} in $\tilde{O}(n^{3/2})$

THEOREM 1.10. *We can certify P_{4k-1} -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

PROOF OF THEOREM 1.10. Let $G = (V, E)$ be an n -vertex graph.

Certification. The certification of the prover is the following. First, it gives to each vertex its certificates in the local computation schemes given by Theorems 4.4 and 4.5, with $\varepsilon = \frac{1}{2}$, so that each vertex u can compute T_G and $G_{\leq u}$. These certificates have size $O(n^{3/2} \log^2 n)$.

Then, the prover adds some additional information in the certificates, in another field, which will be identical in the certificates of all the vertices. We will denote this field by LongestPaths, and it is defined as follows (note that it has size $O(n \log n)$).

Definition 6.1 (P_v , LongestPaths). *For every vertex $v \in V_1$ such that $T_G[v, 2] = (id(C_v), d_v)$, let P_v be a longest induced path starting from v and having all its other vertices in $\bigcup_{0 \leq d < d_v} C_v(d)$, and let $\ell(P_v)$ be its length. We define LongestPaths as being the array indexed by the identifiers of the vertices, such that $\text{LongestPaths}[v] = \ell(P_v)$ for every $v \in \bigcup_{C \in \mathcal{E}_2} \bigcup_{1 \leq d \leq k-1} C(d)$.*

Verification. We define for each $m \geq 1$ a protocol we call m -pathcheck below. (We define the m -pathcheck for arbitrary m , because we will use it again when we prove Theorem 1.13.) In our verification phase each vertex runs the m -pathcheck for $m = 4k - 1$.

Definition 6.2 (m -pathcheck). *Each vertex $u \in V$ takes as input its view at distance k (with the certificates described above) and performs the following steps:*

- (i) *First, u applies the verification of the local computation schemes of Theorems 4.4 and 4.5. If no vertex rejects during this verification phase, then every vertex u computed T_G and $G_{\leq u}$.*
- (ii) *Then, u checks that LongestPaths is the same in its certificate and the certificates of all its neighbors. If it is not the case, u rejects. If no vertex rejects at this point, then LongestPaths is the same in the certificates of all vertices in G .*
- (iii) *Suppose $u \in V_2$. Let C_u denote the ECC_2 containing u . By definition, of ECC and $V_{\leq u}$, $\bigcup_{1 \leq d \leq k-1} C_u(d) \subseteq V_1 \subseteq V_{\leq u}$. Since u computed $G_{\leq u}$ in Step (i), for each $v \in \bigcup_{1 \leq d \leq k-1} C_u(d)$, u can compute the length of the path P_v , and check if it is correctly written in LongestPaths. If this verification fails, u rejects. If no vertex rejects at this point, then LongestPaths is correct in the certificate of every vertex.*

We may assume no vertex has rejected during steps (i) to (iii), and thus every vertex u has access to a (correct) T_G , $G_{\leq u}$ and LongestPaths.

- (iv) If $G_{\leq u}$ has an induced path P on m vertices, such that all the vertices in $P \setminus V_{\leq u}$ are in distinct ECC_2 's, then u rejects.
- (v) If $u \in V_2$, let us denote by C_u its ECC_2 . By definition, $V_{\leq u} = V_1 \cup C_u$ so by step (i) u knows the graph induced by $V_1 \cup C_u$. It rejects if it sees an induced path P_{start} in $G[C_u \cup V_1]$ satisfying all of the following:
 - P_{start} contains at least one vertex from C_u
 - P_{start} ends in a vertex v , such that $T_G[v, 2] = (id(C_v), d_v)$ with $id(C_v) \neq id(C_u)$. This implies that v is at distance $d_v \leq k - 1$ from $C_v \in \mathcal{E}_2 \setminus C_u$. (Recall we assumed that no vertex rejected at step (i) and so u has access to T_G .)
 - v is the unique vertex in P_{start} that is at distance at most d_v from C_v . (again u can check this because it has access to T_G .)
 - $|V(P_{start})| + (\text{LongestPaths}[v]) \geq m$

This case corresponds to the case represented on Figure 5.

Properties of the m -pathcheck. We begin with some observations about the m -pathcheck.

Lemma 6.3. *Suppose G contains an induced path P on m vertices. If there is at most one $C \in \mathcal{E}_2$ such that $|C \cap P| \geq 2$, some vertex rejects during the m -pathcheck.*

PROOF. We may assume that no vertex rejects during the first three steps of the m -pathcheck. If $P \subseteq V_1$, then every vertex $u \in P$ rejects at step (iv). Else, we have $P \cap V_2 \neq \emptyset$, and let $C \in \mathcal{E}_2$ such that $|P \cap C|$ is maximized. By assumption, for every $C' \in \mathcal{E}_2$ such that $C' \neq C$, we have $|P \cap C'| \leq 1$, so every $u \in C$ rejects at step (iv). \square

Lemma 6.4. *Suppose G contains an induced path P on m vertices with the following properties:*

- *there are exactly two distinct ECC_2 's $C_u, C_v \in \mathcal{E}_2$ which contain vertices of P*
- *for some $d \in \{1, 2, \dots, k - 1\}$, there is a unique vertex v at distance exactly d from C_v*

Then some vertex rejects during the m -pathcheck.

PROOF. As depicted in Figure 5, we can decompose P in two consecutive parts P_{start}, P_{end} such that:

- $P_{start} \cap P_{end} = \{v\}$;
- $P_{start} \subseteq C_u \cup V_1$;
- $P_{end} \setminus \{v\} \subseteq \bigcup_{0 \leq d < d_v} C_v(d)$.

We may assume that no vertex rejects during the first four steps of the m -pathcheck. We claim that every $u \in P \cap C_u$ will reject at step (v).

By definition, $|V(P_{start})| = m - |V(P_{end})| - 1$. Let P_v denote the the longest induced path starting at v and having all its other vertices strictly closer from C_v than v . Hence, $|V(P_v)| \geq |V(P_{end})|$, so $\ell(P_{start}) + \text{LongestPaths}[v] \geq m$. P_{start} is an induced path of $G_{\leq u}$, included in $V_{\leq u}$, which satisfies the conditions making u reject at step (v). \square

Correctness. Let us show that this certification scheme is correct. Assume first by contradiction that G contains an induced path P of length $4k - 1$. Let us show that for every assignment of certificates, at least one vertex rejects. If no vertex rejects at step (i) of the verification, then every vertex u knows its witnessed graph $G_{\leq u}$ and the ECC -table T_G (note that u also knows $V_{\leq u}$, because it can be computed directly from T_G). If no vertex rejects in step (iii), then LongestPaths is also correct. Then, let us prove the two following Claims 6.5 and 6.3.




Fig. 5. The decomposition of an induced path in two consecutive paths P_{start} and P_{end} with the properties mentioned in the proof of Theorem 1.10.

Claim 6.5. *There are at most two $C \in \mathcal{E}_2$ such that $|C \cap P| \geq 2$. And if there are exactly two, then P does not contain any vertex from any other ECC_2 .*

PROOF. This simply follows from the fact that P has length $4k - 1$, and two different ECC_2 's are at distance at least $2k - 1$. \square

Applying Claim 6.5 and Lemma 6.3, we can assume that there are exactly two different ECC_2 's which intersect P , and both on at least two vertices. Let us denote them by C_u and C_v . Since P is a path going through C_u and C_v which are at distance at least $2k - 1$ from each other, for all $d \in \{1, \dots, k - 1\}$, we have $P \cap C_u(d) \neq \emptyset$, and $P \cap C_v(d) \neq \emptyset$. Moreover, since P has length $4k - 1$, and since $P \cap C_u(0)$, $P \cap C_v(0)$ both contain at least two vertices, there exists a set $P \cap C_u(d)$ or $P \cap C_v(d)$ which contains only one vertex, for some $d \in \{1, \dots, k - 1\}$. Hence, by applying Lemma 6.4 for $m = 4k - 1$ we obtain that a some vertex rejects during our verification protocol as desired.

Conversely, assume that G does not have an induced path of length $4k - 1$. Let us show that there exists an assignment of the certificates such that every vertex accepts. This assignment is the following one: the prover attributes the certificates such that no vertex reject in the local computation schemes of Theorems 4.4 and 4.5, and gives the correct value of LongestPaths to every vertex. With such a certificate, it is straightforward that a vertex can not reject at steps (i), (ii) and (iii).

Assume by contradiction that a vertex u rejects at step (iv). Then, there exists an induced path P of length $4k - 1$ in $G_{\leq u}$, such that the vertices in $P \setminus V_{\leq u}$ are in distinct ECC_2 's. Since G does not have an induced path of length $4k - 1$, P is not induced in G . It implies that there are two vertices of P linked by an edge in G , and a non-edge in $G_{\leq u}$. Thus, these two vertices are in $V \setminus V_{\leq u}$ (because edges having an endpoint in $V_{\leq u}$ are in $G_{\leq u}$). So by definition of step (iv), these two vertices are in distinct ECC_2 's. In particular, they can not be neighbors in G , which is a contradiction.

Finally, assume by contradiction that a vertex u rejects at step (v). It means that there exists an induced path P_{start} in $G_{\leq u}$, included in $V_{\leq u}$, ending in some vertex v satisfying the conditions of step (v). Since $P_{\text{start}} \subseteq V_{\leq u}$, it is also induced in G . Moreover, there exists an induced path P_v in G

of length $\ell(P_v)$ starting in v and having all its other vertices in $\bigcup_{d < d_v} C_v(d)$. Since v is the only vertex in $P_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$, then the concatenation of P_{start} and P_v is still an induced path in G . Moreover, it has length $\ell(P_{\text{start}}) + \ell(P_v) - 1 \geq 4k - 1$. It is a contradiction. Thus, no vertex rejects at step (v), so all the vertices accept at step (vi), which concludes the proof. \square

6.2 Upper bound for general graphs of size $4k - 1$ in $\tilde{O}(n^{3/2})$

THEOREM 1.11. *For every $k \geq 2$, we can certify H -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$ for every H of size at most $4k - 1$.*

Before proving Theorem 1.11, let us give some informal intuition on the proof. The general idea is the same as in the proof of Theorem 1.10: first, the prover uses the local computation schemes of Theorems 4.4 and 4.5 with $\varepsilon = \frac{1}{2}$ (such that each vertex u computes $G_{\leq u}$ and T_G), and then it adds some additional information in the certificates. The main difference is the following one. Theorem 1.10 is the particular case where the graph H is a path. In this case, the information the prover adds in the certificates is, for each vertex $v \in V$ close to some $C_v \in \mathcal{E}_2$, the length of the longest induced path P_v starting from v and having all its other vertices closer from C_v than v . This enables a vertex u in an $\text{ECC}_2 C_u \neq C_v$ to check, for every induced path P in $G_{\leq u}$, if it can be extended into a path of length $4k - 1$ using P_v . Here, for an arbitrary graph H , we generalize it: the information the prover will add is, intuitively, for each vertex v close from an $\text{ECC}_2 C_v$, and for every induced subgraph H' of H , if H' is induced in the part of G consisting in v and all the vertices closer from C_v than v . With this information, every vertex u in an $\text{ECC}_2 C_u \neq C_v$ will be able to detect if some induced graph in $G_{\leq u}$ can be extended into H using H' .

Pointed graphs. Let us introduce some definitions about pointed graphs that will be useful in the proof of Theorem 1.11. A *pointed graph* is a tuple (H, S) where H is a graph and $S \subseteq V(H)$ is a set of *pointed vertices*. If $(H_1, S_1), (H_2, S_2)$ are two pointed graphs, their *disjoint union* is the pointed graph $(H_1 \cup H_2, S_1 \cup S_2)$, where $H_1 \cup H_2$ is the graph obtained by taking a copy of H_1 and a copy of H_2 , and no edge between them. Finally, let G be a graph, H an induced subgraph of G and $S \subseteq V(H)$. The *complement* of (H, S) in G is the pointed subgraph (\overline{H}, S) where \overline{H} is the subgraph of G induced by $(V(G) \setminus V(H)) \cup S$.

PROOF OF THEOREM 1.11. Let H be a graph with $4k - 1$ vertices. Let us describe a certification scheme for H -free graphs with certificates of size $O(k \cdot 2^{4k} \cdot n + n^{3/2} \log^2 n)$, by looking at distance k .

Certification. The certification of the prover is the following. First, it gives to each vertex its certificates in the local computation schemes given by Theorems 4.4 and 4.5, with $\varepsilon = \frac{1}{2}$. This part of the certificate has size $O(n^{3/2} \log^2 n)$.

Then, the prover adds some additional information in the certificates, in a field which will be identical in the certificates of all the vertices. We will denote this field by HTable , and it is defined as follows. Let $\mathcal{P}(H)$ be the set containing all the pointed graphs $(H', \{h\})$, where $h \in V(H)$, and H' is a subgraph of H induced by $\{h\}$ and an union of connected components of $H \setminus \{h\}$. An example is shown on Figure 6. The field HTable contains a table which has at most n rows and $(4k - 1) \cdot 2^{4k-1}$ columns. The rows are indexed by the identifier of the vertices which are at distance at most $k - 1$ from V_2 , and the columns are indexed by the pointed graphs $(H', \{h\}) \in \mathcal{P}(H)$. Let $v \in V$, such that v is at distance $d_v \leq k - 1$ from some $C_v \in \mathcal{E}_2$. $\text{HTable}[v, (H', \{h\})]$ is equal to 1 if H' is an induced subgraph of $\bigcup_{d \leq d_v} C_v(d)$, where all the vertices in $H' \setminus \{h\}$ are mapped in $\bigcup_{d < d_v} C_v(d)$, and h is mapped to v . Otherwise, $\text{HTable}[v, (H, \{h\})]$ is equal to 0. An example with the pointed graph H' of Figure 6 is shown on Figure 7.




Fig. 6. An example of a graph H and a pointed graph $(H', \{h\}) \in \mathcal{P}(H)$.




Fig. 7. The illustration of $\text{HTable}[v, (H', \{h\})] = 1$, where $(H', \{h\})$ is the pointed graph of Figure 6.

Since the table HTable has at most n rows and $(4k-1) \cdot 2^{4k-1}$ columns, with each entry of constant size, its size is $O(k \cdot 2^{4k} \cdot n)$. The overall size of the certificate is thus $O(k \cdot 2^{4k} \cdot n + n^{3/2} \log^2 n)$.

Verification. The verification of each vertex $u \in V$ consists in the following.

- (i) First, u applies the verification of the local computation schemes of Theorems 4.4 and 4.5. If no vertex rejects during this first verification phase, then every vertex u computed T_G and $G_{\leq u}$.
- (ii) Then, u checks that HTable is the same in its certificate and in the certificate of all its neighbors. If it is not the case, u rejects. If no vertex rejects, HTable is the same in the certificates of all the vertices in G . The next step will consist in checking its correctness.
- (iii) If $u \in V_2$, it does the following. Let C_u be its ECC_2 . We have $\bigcup_{0 \leq d \leq k-1} C_u(d) \subseteq V_{\leq u}$. Thus, for every $v \in \bigcup_{0 \leq d \leq k-1} C_u(d)$ and every $(H, \{h\}) \in \mathcal{P}(H)$, u can check if $\text{HTable}[v, (H', \{h\})]$ is correct. If it is not correct, u rejects. If no vertex rejects at this point, then HTable is correct.
- (iv) If $G_{\leq u}$ has an induced H , such that all the vertices in $H \setminus V_{\leq u}$ are in distinct ECC_2 's, then u rejects (in particular, this is the case if $H \subseteq V_{\leq u}$). The last steps are similar to step (v) in the verification of Theorem 1.10: u will check that it can not obtain a graph H by extending a subgraph of $G_{\leq u}$ into H , using a subgraph H' such that $\text{HTable}[v, (H', \{h\})] = 1$ for some vertices v and h .
- (v) If $u \in V_2$, let $C_u \in \mathcal{E}_2$ such that $u \in C_u$. The vertex u rejects if there exists two pointed graphs $(H_{\text{start}}, \{h\}), (H_{\text{end}}, \{h\}) \in \mathcal{P}(H)$ which are complementary of each other in H , and a vertex $v \in C_v(d_v)$ for some $C_v \neq C_u, d_v \leq k-1$, such that the following conditions are satisfied:




Fig. 8. Illustration of the condition making u reject at step (v) of the verification. Here, the graph H is the one of Figure 6. The pointed graph $(H_{\text{end}}, \{h\})$ is the pointed graph $(H', \{h\})$ of Figure 6. .

- H_{start} is an induced subgraph of $G_{\leq u}$, included in $V_{\leq u}$, which contains at least one vertex from C_u
- h is mapped to v , and it is the only vertex in $H_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$
- $\text{HTable}[v, (H_{\text{end}}, \{h\})] = 1$, where $(H_{\text{end}}, \{h\})$ is the complementary of $(H_{\text{start}}, \{h\})$ in H

An illustration is shown on Figure 8.

(vi) If $u \in V_2$, let $C_u \in \mathcal{E}_2$ such that $c \in C_u$. The vertex u rejects if there exists three pointed graphs $(H_{\text{start}}, \{h_1, h_2\})$, $(H_{\text{end}}^{(1)}, \{h_1\})$, $(H_{\text{end}}^{(2)}, \{h_2\})$, and two vertices $v_1 \in C_{v_1}(d_{v_1})$, $v_2 \in C_{v_2}(d_{v_2})$ for some $d_{v_1}, d_{v_2} \leq k - 1$ where C_{v_1}, C_{v_2} are two different ECC_2 's different from C_u , such that the following conditions are satisfied:

- $H_{\text{end}}^{(1)}, H_{\text{end}}^{(2)}$ are two disjoint induced subgraphs of H
- $(H_{\text{start}}, \{h_1, h_2\})$ is the complementary in H of the disjoint union of $(H_{\text{end}}^{(1)}, \{h_1\})$ and $(H_{\text{end}}^{(2)}, \{h_2\})$
- H_{start} is an induced subgraph of $G_{\leq u}$, included in $V_{\leq u}$, which contains at least vertex from C_u
- h_1 (resp. h_2) is mapped to v_1 (resp. v_2), and it is the only vertex in $H_{\text{start}} \cap (\bigcup_{d \leq d_{v_1}} C_{v_1}(d))$ (resp. in $H_{\text{start}} \cap (\bigcup_{d \leq d_{v_2}} C_{v_2}(d))$)
- $\text{HTable}[v_1, (H_{\text{end}}^{(1)})] = 1$ and $\text{HTable}[v_2, (H_{\text{end}}^{(2)})] = 1$

An illustration is shown on Figure 9.

(vii) If u did not reject at this point, it accepts.

Correctness. Let us show that this certification scheme is correct. First, assume that G contains an induced H , and let us show that for every assignment of certificates, at least one vertex rejects. If no vertex rejects at step (i) of the verification, then every vertex u computed the ECC -table T_G and its witnessed graph $G_{\leq u}$ (u knows also $V_{\leq u}$ because it can be computed directly from T_G). If no vertex rejects at step (iii), then HTable is also correct. Then, we have the following Claim 6.6, and its proof is analogous as proofs of Claims 6.5 and 6.3.

Claim 6.6. (1) *There are at most three $C \in \mathcal{E}_2$ such that $C \cap H \neq \emptyset$.*

(2) *If there is at most one $C \in \mathcal{E}_2$ such that $|C \cap H| \geq 2$, then there exists a vertex which rejects at step (iv) of the verification.*

By Claim 6.6, we can assume that there are at least two $C \in \mathcal{E}_2$ such that $|C \cap H| \geq 2$. Now, there are two cases.




Fig. 9. Illustration of the condition making u reject at step (vi) of the verification. The graph H is the one of Figure 6. The graph $H_{\text{end}}^{(1)}$ is a path on two vertices and the graph $H_{\text{end}}^{(2)}$ is a triangle.

- If there are exactly two ECC_2 's intersecting H , denoted by C_1 and C_2 , since H has $4k - 1$ vertices and that, by assumption, we have $|C_i \cap H| \geq 2$ for $i \in \{1, 2\}$, there exists $1 \leq d \leq k - 1$ and $i \in \{1, 2\}$ such that $|C_i(d)| = 1$. Without loss of generality, assume that $i = 1$, and let us denote by v the unique vertex in $C_1(d)$. Let $u \in C_2$. Then, we can decompose H in two parts H_{start} and H_{end} satisfying the conditions making u reject at step (v).
- If there are exactly three ECC_2 's intersecting H , denoted by C_1 , C_2 and C_3 , there is at most one $i \in \{1, 2, 3\}$ such that, for all $d \in \{1, \dots, k - 1\}$, $|C_i(d)| \geq 2$. Without loss of generality, assume that there exists $d_1, d_2 \in \{1, \dots, k - 1\}$ such that $|C_i(d_i)| = 1$ for all $i \in \{1, 2\}$. Let us denote by v_1 (resp. by v_2) the only vertex in $C_1(d_1)$ (resp. in $C_2(d_2)$). Let $u \in C_3$. Then, we can decompose H in three parts $H_{\text{start}}, H_{\text{end}}^{(1)}, H_{\text{end}}^{(2)}$ satisfying the conditions making u reject at step (vi).

Conversely, assume that G does not have an induced H , and let us show that there exists an assignment of the certificates such that every vertex accepts. This assignment is the following one: the prover attributes the certificates such that no vertex rejects in the local computation schemes of Theorems 4.4 and 4.5, and gives the correct value of H_{Table} to every vertex. With such a certificate, a vertex can not reject at steps (i), (ii) and (iii). With exactly the same argument as in the proof of Theorem 1.11, a vertex can not reject at step (iv).

By contradiction, assume that a vertex u rejects at step (v). Then, there exists two pointed graphs $(H_{\text{start}}, \{h\})$, $(H_{\text{end}}, \{h\})$ satisfying the conditions making u reject at step (v). By the properties they satisfy, we can glue them in h to obtain a graph H , which is still induced in G , which is a contradiction. Similarly, if a vertex u rejects at step (vi), there exists three pointed graphs $(H_{\text{start}}, \{h_1, h_2\})$, $(H_{\text{end}}^{(1)}, h_1)$, $(H_{\text{end}}^{(2)}, h_2)$ satisfying the conditions making u reject at step (vi). We can again glue them on h_1, h_2 to obtain an induced H in G , which is a contradiction. Thus, all the vertices accept at step (vii), which concludes the proof. \square

We stated Theorem 1.11 with an induced subgraph perspective but one can easily remark that the proof technique is general enough to be directly extended to subgraphs rather than

induced subgraphs. Indeed, we simply have to change the definition of HTable in such a way $\text{HTable}[v, (H', \{h\})]$ is equal to 1 if H' is a subgraph of $\bigcup_{d \leq d_v} C_v(d)$ instead of an induced subgraph in order to get the same result. We actually think that our proof technique can be used more generally to get subquadratic bounds in the polynomial regime.

6.3 Upper bound for paths of length at most $3k - 1$ in $\tilde{O}(n)$

THEOREM 1.12. *For every $k \geq 2$, we can certify P_{3k-1} -free graphs by looking at distance k with certificates of size $O(n \log^3 n)$.*

To prove this result, we will in fact prove the following Proposition 6.7, and Theorem 1.12 will be a simple corollary. While, until now, all the results were based on $\varepsilon = \frac{1}{2}$ and the partition of the vertex set into two parts, we will use the whole power of our machinery in this proof by choosing an arbitrarily small value of ε ; this leads to new technicalities we have to deal with in the proof.

Proposition 6.7. *There exists $c > 0$ such that, for all $0 < \varepsilon < 1$, there exists a certification scheme for P_{3k-1} -free graphs using at most $\frac{c}{\varepsilon} \cdot \log^2 n \cdot n^{1+\varepsilon}$ bits.*

PROOF OF THEOREM 1.12 ASSUMING PROPOSITION 6.7. The certification in size $O(n \log^3 n)$ is the following one. Let n be the number of vertices in the graph. The prover writes n in the certificate of every vertex of the graph, and certifies the correctness of this information with $O(n \log n)$ bits (by coding a spanning tree). Then, it uses the certification scheme given by Proposition 6.7, with $\varepsilon = 1/\log n$. This uses $O(n \log^3 n)$ bits in total, since $n^{1/\log n}$ is a constant. \square

PROOF OF PROPOSITION 6.7. Let $0 < \varepsilon < 1$ and $N = \lceil \frac{1}{\varepsilon} \rceil$. Let us describe a certification scheme for P_{3k-1} -free graphs. We will show that it is correct and uses at most $\frac{c}{\varepsilon} \cdot \log^2 n \cdot n^{1+\varepsilon}$ bits, where c is a constant independent of ε . Let $G = (V, E)$ be a graph and $n = |V|$.

Certification. Let us describe the certificates given by the prover to the vertices. First, the prover gives to every vertex its certificates in the local computation schemes of Theorems 4.4 and 4.5, so that each vertex u can compute T_G and $G_{\leq u}$. These certificates have size $O(\frac{n^{1+\varepsilon}}{\varepsilon} \log^2 n)$. Then, as in the certification scheme of Theorem 1.10, the prover adds some information in the certificates, in another field, denoted by LongestPaths, which is identical in the certificates of all the vertices. In LongestPaths, the prover writes a table which has n rows and N columns. The rows are indexed with the identifiers of the vertices, and the columns by $\{1, \dots, N\}$. Let $v \in V$ and $i \in \{1, \dots, N\}$. If $v \notin \bigcup_{C \in \mathcal{E}_i} \bigcup_{0 \leq d \leq k-1} C(d)$, we set $\text{LongestPaths}[v, i] := \perp$. Else, there exists a unique $C_v \in \mathcal{E}_i$ and $d_v \leq k - 1$ such that $v \in C_v(d_v)$. In this case, let P_v denote the longest induced path starting from v and having all its other vertices in $\bigcup_{0 \leq d < d_v} C_v(d)$, and let us denote its length by $\ell(P_v)$. We set $\text{LongestPaths}[v, i] := \ell(P_v)$.

Since LongestPaths has n rows, N columns, and each entry is written on at most $\log n$ bits, its size is $O(Nn \log n)$. Thus, the overall size of each certificate remains $O(\frac{n^{1+\varepsilon}}{\varepsilon} \log^2 n)$.

Verification. The verification of each vertex $u \in V$ consists in the following steps:

- (i) First, u applies the verification of the local computation schemes given by Theorems 4.4 and 4.5.

If no vertex rejects, then every vertex u computed its witnessed graph $G_{\leq u}$ and the ECC-table T_G of G .

- (ii) Then, u checks that the field LongestPaths is written the same in its certificate and the certificates of all its neighbors. If it is not the case, u rejects.

If no vertex rejects, LongestPaths is the same in the certificates of all the vertices in G . The next step will consist in checking its correctness.




Fig. 10. The decomposition of P in three consecutive parts P_1, P_2, P_3 . For each $i \in \{1, 2, 3\}$, x_i is the vertex of maximum degree in P_i .

- (iii) Let $i_u \in \{1, \dots, N\}$ be such that $u \in V_{i_u}$. For each $i \in \{1, \dots, i_u\}$, u does the following. We have $u \in H_i$. Let us denote by C_u its ECC $_i$. Let $d \in \{1, \dots, k-1\}$, and let $v \in C_u(d)$. Since $v \notin C_u$, we have $v \in L_{i-1}$ (indeed, if $v \in H_i$, it would imply $v \in C_u$ because v is at distance at most $k-1$ from C_u). Let $i_v \in \{1, \dots, i-1\}$ be such that $v \in V_{i_v}$. We have $C_u \subseteq H_i \subseteq H_{i_v}$, so u, v are in the same ECC $_{i_v}$. Thus, we have $v \preceq u$. Since this true for all $d \in \{1, \dots, k-1\}$ and $v \in C_u(d)$, we have $\bigcup_{1 \leq d \leq k-1} C_u(d) \subseteq V_{\preceq u}$. Thus, for all $v \in \bigcup_{1 \leq d \leq k-1} C_u(d)$, u can check if $\text{LongestPaths}[v, i]$ is correct, and rejects if it is not the case.
- If no vertex rejects at this point, then LongestPaths is correct.
- (iv) If $G_{\preceq u}$ has an induced path P of length $3k-1$, which is included in $V_{\preceq u}$, then u rejects.
- (v) Finally, u rejects if it sees, in its view at distance k , an induced path P of some length $\ell(P)$, which goes from a vertex v_1 to a vertex v_2 , such that the following conditions are satisfied:
- at most one vertex in P is at distance exactly k from u
 - there exists $i \in \{1, \dots, N\}$ and two distinct $C_{v_1}, C_{v_2} \in \mathcal{E}_i$ and $d_{v_1}, d_{v_2} \in \{0, \dots, k-1\}$ such that $v_1 \in C_{v_1}(d_{v_1})$ and $v_2 \in C_{v_2}(d_{v_2})$
 - v_1 (resp. v_2) is the only vertex in $P \cap (\bigcup_{0 \leq d \leq d_{v_1}} C_{v_1}(d))$ (resp. in $P \cap (\bigcup_{0 \leq d \leq d_{v_2}} C_{v_2}(d))$)
 - $\ell(P) + (\text{LongestPaths}[v_1, i] - 1) + (\text{LongestPaths}[v_2, i] - 1) \geq 3k-1$
- (vi) If u did not reject previously, it accepts.

Correctness. Let us show that this certification scheme is correct. First, assume that G has an induced path P of length $3k-1$, and let us show that for every assignment of certificates, at least one vertex rejects. If no vertex rejects at step (i), then every vertex u knows the ECC-table T_G , and its witnessed graph $G_{\preceq u}$. If no vertex rejects in steps (ii) and (iii), then LongestPaths is correctly written in the certificates.

Let us decompose P in three consecutive parts P_1, P_2, P_3 , having respectively $k, k-1, k$ vertices. Let x_1 (resp. x_2, x_3) be the vertex of maximum degree in P_1 (resp. in P_2, P_3). Let $i_1, i_2, i_3 \in \{1, \dots, N\}$ be such that $x_1 \in V_{i_1}, x_2 \in V_{i_2}$ and $x_3 \in V_{i_3}$ (see Figure 10 for an illustration). Note that the distance between x_1 and x_2 is at most $2k-2$. Similarly, the distance between x_2 and x_3 is at most $2k-2$. The two following Claims 6.8 and 6.9 show that, in all cases, at least one vertex rejects.

Claim 6.8. *If $i_2 \geq \min(i_1, i_3)$, then at least one vertex reject at step (iv) of the verification.*

PROOF. Assume that $i_2 \geq i_1$. Then, we have $x_2 \in H_{i_1}$. Since x_1 and x_2 are at distance at most $2k-2$, there are in the same ECC $_{i_1}$. Thus, we have $x_1 \preceq x_2$. Note also that, since x_1 is the vertex of maximum degree in P_1 , we have $P_1 \subseteq V_{\preceq x_1}$. Thus, by transitivity, we have $P_1 \subseteq V_{\preceq x_2}$. Moreover, we also have $P_2 \subseteq V_{\preceq x_2}$. Finally, there are two cases:

- If $i_2 \geq i_3$, by the same argument, we have $P_3 \subseteq V_{\preceq x_2}$. So P is included in $V_{\preceq x_2}$, and x_2 rejects at step (iv).
- If $i_3 > i_2$, we have $x_2 \preceq x_3$. By transitivity we have $P_1 \subseteq V_{\preceq x_3}$ and $P_2 \subseteq V_{\preceq x_3}$. Since we also have $P_3 \subseteq V_{\preceq x_3}$, P is included in $V_{\preceq x_3}$ and x_3 rejects at step (iv). \square

Claim 6.9. *If $i_2 \leq \min(i_1, i_3)$, then at least one vertex rejects.*




Fig. 11. The vertex u_1 (resp. u_3) is the only vertex in $P \cap C_{x_1}(d_1)$ (resp. in $P \cap C_{x_3}(d_3)$). The vertices u_1 and u_3 are at distance at most $2k - 1$ from each other. The vertex u is such that its distance to u_1 is at most k , and its distance to u_3 is at most $k - 1$.

PROOF. Without loss of generality, assume that $i_2 \leq i_3 \leq i_1$. We have $x_1 \in H_{i_3}$. There are two cases, depending on whether x_1 and x_3 are in the same ECC_{i_3} or not.

- If x_1 and x_3 are in the same ECC_{i_3} , we have $x_3 \leq x_1$. Thus, by transitivity, $P_3 \subseteq V_{\leq x_1}$. Since we also have $P_1 \subseteq P_{\leq x_1}$ and $P_2 \in V_{\leq x_1}$ (because $i_2 \leq i_1$), then P is included in $V_{\leq x_1}$ so x_1 rejects at step (iv).
- If x_1 and x_3 are not in the same ECC_{i_3} , let us prove that some vertex rejects at step (v). Let $C_{x_1}, C_{x_3} \in \mathcal{E}_{i_3}$ be such that $x_1 \in C_{x_1}$ and $x_3 \in C_{x_3}$. Since P passes through C_{x_1} and C_{x_3} , in the part between x_1 and x_3 , it passes also through $C_{x_1}(d)$ and $C_{x_3}(d)$ for all $d \in \{0, \dots, k-1\}$. Moreover, since P has length $3k-1$, there exists at least $k+1$ sets among $\{C_{x_1}(0), \dots, C_{x_1}(k-1), C_{x_3}(0), \dots, C_{x_3}(k-1)\}$ through which P passes only once. Thus, in the part of P between x_1 and x_3 , there exists $d_1, d_3 \in \{0, \dots, k-1\}$ such that P passes only once through $C_{x_1}(d_1)$ and $C_{x_3}(d_3)$, in some vertices denoted by u_1, u_3 , and u_1, u_3 are at distance at most $2k-1$ from each other. Let P' denote the part of P between u_1 and u_3 . Let $u \in P'$ such that u is at distance at most $k-1$ from every vertex in P' , except u_1 , and u, u_1 are at distance at most k . This is depicted on Figure 11.

So P' is an induced path in the view of u at distance k . Let us denote its length by $\ell(P')$. Since P has length $3k-1$, by definition of LongestPaths, we have:

$$(\text{LongestPaths}[u_1, i_3] - 1) + \ell(P') + (\text{LongestPaths}[u_3, i_3] - 1) \geq 3k - 1$$

Thus, the vertex u rejects at step (vi). \square

Conversely, assume that there is no induced P_{3k-1} in G . Then, let us prove that with the certificates described above, no vertex rejects. It is straightforward that no vertex rejects at steps (i), (ii) and (iii). By contradiction, assume that a vertex u rejects at step (iv). Then, there exists an induced path P of length $3k-1$ in $G_{\leq u}$ which is included in $V_{\leq u}$. So P is also induced in G , which is a contradiction. Finally, assume that a vertex u rejects at step (v). Then, there is a path P induced in the view of u at distance k and two vertices v_1, v_2 satisfying the conditions making u reject at step (v). Let $i \in \{1, \dots, N\}$ be such that v_1 and v_2 are at distance at most $k-1$ from two distinct ECC_i 's, denoted by C_{v_1} and C_{v_2} . Let d_{v_1} (resp. d_{v_2}) be such that $v_1 \in C_{v_1}(d_{v_1})$ (resp. $v_2 \in C_{v_2}(d_{v_2})$). By definition of LongestPaths, there exists a path P_{v_1} (resp. P_{v_2}) which starts in v_1 (resp. v_2), which has all its other vertices in $\bigcup_{0 \leq d < d_{v_1}} C_{v_1}(d)$ (resp. in $\bigcup_{0 \leq d < d_{v_2}} C_{v_2}(d)$), and such that $\ell(P) + (\ell(P_{v_1}) - 1) + (\ell(P_{v_2}) - 1) \geq 3k - 1$. So we can glue the paths P_{v_1}, P and P_{v_2} to obtain a path of length $3k-1$ in G , which is still induced. This is a contradiction. Thus, all the vertices accept at step (vi). \square

6.4 Upper bound for paths of length $\frac{14}{3}k - 2$ in $\tilde{O}(n^{3/2})$

In this subsection we prove the following theorem.

THEOREM 1.13. *For every $k \geq 2$, we can certify $P_{\lceil \frac{14}{3}k \rceil - 1}$ -free graphs by looking at distance k with certificates of size $O(n^{3/2} \log^2 n)$.*

PROOF. Let $m = \lceil \frac{14}{3}k \rceil - 1$. Note that $m < \frac{14}{3}k$. Let us describe a certification scheme for P_m -free graphs, with certificates of size $O(n^{3/2} \log^2 n)$. Let $G = (V, E)$ be a graph and $n = |V|$. Let $\varepsilon = \frac{1}{2}$. We will use the definitions of Section 4.2.1 with respect to ε . We remind the reader of the following key pieces of notation.

- V_1 is the set of vertices with degree less than \sqrt{n} . $V_2 = V \setminus V_1$.
- $H_1 = V$, $H_2 = V_2$.
- \mathcal{E}_2 is the partition of H_2 into ECC_2 's. By definition all of V is in one ECC_1 .
- For $C \in \mathcal{E}_2$, we denote by $C(d)$ the set of vertices at distance exactly d from C .
- If u is at distance at most $k - 1$ from some ECC_2 C_u , $T_G[u, 2]$ is $(id(C_u), d_u)$ where d_u is the distance from u to C_u . (Recall, C_u is uniquely defined.) Otherwise, $T_G[u, 2] = \perp$.
- For each $u \in V_1$, by definition $V_{\leq u} = V_1$. For each $u \in V_2$, by definition $V_{\leq u} = V_1 \cup C_u$ where C_u denotes the ECC_2 containing u .
- $G_{\leq u}$ is the graph obtained from G by deleting all edges that do not have at least one endpoint in $V_{\leq u}$.
- For every vertex $v \in V_1$ such that v is at distance at most $k - 1$ from a (unique) ECC_2 C_v , we let P_v be the longest induced path starting from v with other vertices in P_v strictly closer to C_v than v . Let $\ell(P_v)$ be the length of P .
- We define `LongestPaths` to be an array indexed by identifiers of vertices, such that `LongestPaths[v] = $\ell(P_v)$` for every vertex $v \in \bigcup_{C \in \mathcal{E}_2} \bigcup_{1 \leq d \leq k-1} C(d)$, and it is \perp otherwise.

Certification. Let us describe the certificates given by the prover to the vertices.

We give every vertex v the certificates of Theorem 1.10 and some additional information denoted by `LongestConstrainedPath`. We repeat the definition of the certificates of Theorem 1.10 for the readers convenience: First, it gives to every vertex its certificates in the local computation schemes of Theorems 4.4 and 4.5, with ε , so that each vertex u can compute T_G and $G_{\leq u}$. This part of the certificate has size $O(n^{3/2} \log^2 n)$. Then the prover adds the field `LongestPaths` to each vertex. It has size $O(n \log n)$.

For every $v \in V$, let us describe what information is stored in `LongestConstrainedPath(v)`. If $v \in V_2$ or $v \in V_1$ and v has distance at least k from every ECC_2 , then we set `LongestConstrainedPath(v) = \perp` . Otherwise, `LongestConstrainedPath(v)` consists in a table having at most n rows, and 4 columns. The rows are indexed by vertices of G . The content of each cell is the following. Let $v' \in C_v(d_v) \setminus \{v\}$. Let Q denote the set $\bigcup_{0 \leq d < d_v} C_v(d)$. In other words, Q is the set of vertices in G with that are strictly closer to C_v than v .

- In `LongestConstrainedPath(v)[v', 1]`, the prover writes the length of the longest induced path starting from v and having all the other vertices in $Q \cup \{v'\}$.
- In `LongestConstrainedPath(v)[v', 2]`, the prover writes the length of the longest induced path starting from v' and having all the other vertices in $Q \setminus N[v]$.
- In `LongestConstrainedPath(v)[v', 3]`, the prover writes the length of the longest induced path starting at v , ending in v' , and having all the other vertices in Q .
- In `LongestConstrainedPath(v)[v', 3]`, the prover writes the maximum sum of the lengths of two disjoint induced paths, which are anti-complete one to each other, which start respectively at v and v' , and which have all the other vertices in Q .

Note that `LongestPaths` and `LongestConstrainedPath` both have size $O(n \log n)$. Thus, the overall size of each certificate is $O(n^{3/2} \log^2 n)$.

Verification. The verification of each vertex $u \in V$ consists in the following steps. First u performs the m -pathcheck as described in Definition 6.2. (Recall $m = \lceil 14k/3 \rceil - 1$.) We will refer to the steps of Definition 6.2 as step (i) - step (v).

Then u performs the following additional steps:

(vi) u checks the correctness of $\text{LongestConstrainedPath}(u)$.

- Since V_1, V_2 are defined only by the degrees of vertices, u knows whether it is in V_1 or V_2 . If $u \in V_2$, and $T_G[u, 2] \neq \perp$, u rejects. If $u \in V_1$, u checks whether $\text{LongestConstrainedPath}(u) = \perp$ if and only if $T_G[u, 2] = \perp$. If not, u rejects.
- Suppose $\text{LongestConstrainedPath}(u) \neq \perp$. Then by the previous item and the definition of T_u , there is a (unique) $C_u \in \mathcal{E}_2$ and $d_u \in \{1, \dots, k-1\}$ such that u is at distance exactly d_u from C_u . Let $u' \in C_u$ at distance d_u from u . By Remark 4.7, u can compute $G_{\leq u'}$. In particular, u knows the subgraph of G induced by $\bigcup_{0 \leq d \leq k-1} C_u(d)$. and this is sufficient for u to check the correctness of $\text{LongestConstrainedPath}(u)$ by definition of $\text{LongestConstrainedPath}$ (and u rejects if it is not correct.)

We may assume no vertex has rejected at this point, and thus u has access to a (correct) T_G , $G_{\leq u}$, LongestPaths and $\text{LongestConstrainedPath}(u)$.

(vii) If $u \in V_2$, let us denote by C_u its ECC_2 . It rejects if it sees an induced path P_{start} in $G[C_u \cup V_1]$, and such that the following conditions are satisfied:

- P_{start} contains at least one vertex from C_u
- P_{start} goes from a vertex v to a vertex w , such that $T_G[v, 2] = (id(C_v), d_v)$, $T_G[w, 2] = (id(C_w), d_w)$, and $\text{Id}(C_u), \text{Id}(C_v), \text{Id}(C_w)$ are all distinct. Since T_G is correct by step (i), C_u, C_v, C_w are all distinct, v is at distance at most $d_v \leq k-1$ from C_v and w is at distance at most $d_w \leq k-1$ from C_w .
- v is the only vertex in $P_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$, and w is the only vertex in $P_{\text{start}} \cap (\bigcup_{d \leq d_w} C_w(d))$. In other words all other vertices of P_{start} are further from C_v (resp. C_w) than v (resp. w). u checks this by looking at the entries of T_G for the vertices in P .
- $\ell(P_{\text{start}}) + (\text{LongestPaths}[v] - 1) + (\text{LongestPaths}[w] - 1) \geq m$

This case is represented on Figure 12.

(viii) Finally, if $u \in V_1$ and u is at distance $d_u \in \{1, \dots, k-1\}$ from some vertex u' in some ECC_2 C_u we perform the following (recall u can check this condition because u has access to T_G). The vertex u computes $G_{\leq u'}$ (this is possible by Remark 4.7), so u knows the graph $G[C_u \cup V_1]$. Then, u rejects if there exists an induced path P_{start} in $G[C_u \cup V_1]$, $C_v \in \mathcal{E}_2 \setminus \{C_u\}$, $d_v \in \{1, \dots, k-1\}$ and two vertices $v, v' \in C_v(d_v)$, with v at distance at most k from u , such that any of the following set of conditions is satisfied (see Figure 13 for an illustration):

- (a)
 - P_{start} ends in v and is anti-complete to v'
 - v is the only vertex in $P_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$
 - $\ell(P_{\text{start}}) + (\text{LongestConstrainedPath}(v)[v', 1] - 1) \geq m$
- (b)
 - P_{start} ends in v' and passes through v
 - v, v' are the only vertices in $P_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$
 - $\ell(P_{\text{start}}) + (\text{LongestConstrainedPath}(v)[v', 2] - 1) \geq m$
- (c)
 - P_{start} ends in v , and there exists an induced path in $G[C_u \cup V_1]$, denoted by P_{end} , which starts in v' and is anti-complete to P_{start}
 - v, v' are the only vertices in $(P_{\text{start}} \cup P_{\text{end}}) \cap (\bigcup_{d \leq d_v} C_v(d))$
 - $\ell(P_{\text{start}}) + (\text{LongestConstrainedPath}(v)[v', 3] - 2) + \ell(P_{\text{end}}) \geq m$
- (d)
 - P_{start} starts in v and ends in v'
 - v, v' are the only vertices in $P_{\text{start}} \cap (\bigcup_{d \leq d_v} C_v(d))$




Fig. 12. The case making u reject at step (vii) in the certification scheme of Theorem 1.13

- $\ell(P_{\text{start}}) + (\text{LongestConstrainedPath}(v)[v', 4] - 2) \geq m$




Fig. 13. The four cases making u reject at step (viii) of the verification algorithm in Theorem 1.13. In each of these cases, the vertices u and v are at distance at most k (which enables u to see the certificate of v).

(ix) If u did not reject in the previous steps, it accepts.

Correctness. Let us prove that this certification scheme is correct. First let us assume for a contradiction, that G contains an induced path P of length m and that every vertex accepts. Using the m -pathcheck, every vertex u can correctly compute T_G , $G_{\leq u}$, and knows that the values of

LongestPaths and LongestConstrainedPath(u) written by the prover in the certificates are correct. Then, let us prove the following Claims 6.10, 6.12 and 6.13.

Claim 6.10. *There are three subpaths P_1, P_2, P_3 of P such that $P = P_1 \cup P_2 \cup P_3$ and for each $i \in \{1, 2, 3\}$ there is a $C_i \in \mathcal{E}_2$ such that every vertex of $V(P_i) \subseteq C_i \cup V_1$.*

PROOF. This simply follows from the fact that P has length $m < \frac{14}{3}k$ and that two distinct ECC_2 's are at distance at least $2k$ from each other by definition of extended connected component. ■

It follows from Claim 6.10 that there are at most three distinct $C \in \mathcal{E}_2$ such that $P \cap C \neq \emptyset$.

Claim 6.11. $P \cap V_2 \neq \emptyset$

PROOF. Recall $V_1 \subseteq V_{\leq u}$ for every vertex $u \in V$. In step (iv) of the m -pathcheck, every vertex checks whether there is an induced path of length m with vertex set in $V_{\leq u}$. ■

By applying Lemma 6.3 for m , we obtain that:

Claim 6.12. *There are at least two distinct $C \in \mathcal{E}_2$ such that $|P \cap C| \geq 2$.*

Claim 6.13. *If there are exactly three $C \in \mathcal{E}_2$ such that $P \cap C \neq \emptyset$, at least one vertex rejects.*

PROOF. Assume that there are exactly three $C \in \mathcal{E}_2$ such that $P \cap C \neq \emptyset$. Let us denote these three ECC_2 's by C_1, C_2, C_3 . By Claim 6.12, we can also assume that P has at least 5 vertices in $C_1 \cup C_2 \cup C_3$. Without loss of generality, there is a subpath of P passing through C_1, C_3, C_2 , in that order. For $x \in \{1, 2, 3\}$, let B_x denote the set of vertices of $V(P)$ in $\bigcup_{1 \leq d \leq k-1} C_x(d)$. Since any two distinct ECC_2 's are at distance at least $2k-1$ from each other, $|B_1|, |B_2| \geq k-1$ and $|B_3| \geq 2k-2$. By definition of ECC , B_1, B_2, B_3 are pairwise disjoint. By Claim 6.12, we can also assume that P has at least 5 vertices in $C_1 \cup C_2 \cup C_3$. And by definition, $C_1 \cup C_2 \cup C_3$ is disjoint from $B_1 \cup B_2 \cup B_3$. By definition, $|V(P)| = m \leq \frac{14}{3}k$. So by combining these observations we obtain:

$$|B_1|, |B_2| \leq V(P) - 3k + 3 - 5 < 2k - 2 \quad (1)$$

Let $i \in \{1, 2\}$. From the definition of ECC , there must be a vertex at distance at $k-1$ from C_i in B_i . Since P contains a vertex of C_i it follows that for each $d \in \{1, 2, \dots, k-1\}$ there is some vertex $b \in B_i$ at distance exactly d from C_i . Hence, it follows from (1) that there exists some $d_i \in \{1, \dots, k-1\}$ such that $|P \cap C_i(d_i)| = 1$. Let v_i denote the unique element of $P \cap C_i(d_i)$. Let P_{start} denote the $v_1 v_2$ -path of P .

Let $u \in C_3$ be arbitrary. We will show that u rejects at step (vii). By definition of ECC_2 , the only vertices of V_2 in P_{start} are contained in C_3 . So in particular, $V(P_{\text{start}}) \subseteq P_{\leq u}$. Hence, u will consider P_{start} in step (vii) and it will reject if $\ell(P_{\text{start}}) + (\text{LongestPaths}[v_1] - 1) + (\text{LongestPaths}[v_2] - 1) \geq m$.

By definition, there the graph obtained from P by deleting all vertices of P_{start} except for v_1, v_2 consists of two paths P_1, P_2 such that v_1 is an end of P_1 and v_2 is an end of P_2 . Let $i \in \{1, 2\}$. By Claim 6.10 we have that $P_i \subseteq V_i \cup C_i$. Since P_i contains a vertex of C_i , by the definition of distance, there can be no vertex $V(P_i)$ at distance strictly greater than d_i from v_i . It follows that $V(P_i) \setminus \{v_i\} \subseteq \bigcup_{0 \leq d < d_i} C_i(d)$. Hence by definition, $\text{LongestPaths}(v_i) \geq |V(P_i)| - 1$.

Since $P = P_1 \cup P_2 \cup P_3$, it follows that $\ell(P_{\text{start}}) + (\text{LongestPaths}[v_1] - 1) + (\text{LongestPaths}[v_2] - 1) \geq m$ and therefore u rejects in step (vii). □

Applying Claims 6.10, 6.12 and 6.13, we can assume that there exactly two ECC_2 's which intersect P , both on at least two vertices. Let us denote these two ECC_2 's by C_u and C_v .

Claim 6.14. *For all $d \in \{1, \dots, k-1\}$, we have $|P \cap C_v(d)| \geq 2$, and similarly $|P \cap C_u(d)| \geq 2$.*

PROOF. Since C_1, C_2 both meet $V(P)$, by definition of ECC for each $C \in \{C_v, C_u\}$ there is at least one vertex in P at distance exactly d from C . If there exists $d \in \{1, \dots, k-1\}$ and $x \in \{u, v\}$ such that $|P \cap C_x(d)| = 1$, then some vertex will reject by Lemma 6.4. ■

For brevity, let $\mathcal{B}_u := \bigcup_{0 \leq d \leq k-1} C_u(d)$ and let $\mathcal{B}_v := \bigcup_{0 \leq d \leq k-1} C_v(d)$. Let us denote by α the number of vertices in $P \setminus (\mathcal{B}_u \cup \mathcal{B}_v)$. In other words, α denotes the number of vertices in P that are not “close” (at most $k-1$ away from) C_u or C_v . Thus by Claim 6.14,

$$\alpha \leq m - 4(k-1) \leq \frac{14}{3}k - 4k - 4 < \frac{2}{3}k. \quad (2)$$

Recall, that by definition \mathcal{B}_u and \mathcal{B}_v are disjoint. So, $|V(P)| = \lceil 14k/3 \rceil - 1 = \alpha + |P \cap \mathcal{B}_u| + |P \cap \mathcal{B}_v|$. Hence, by the pigeonhole principle and the symmetry between C_u and C_v we may assume that

$$|P \cap \mathcal{B}_v| < \frac{7}{3}k - \frac{\alpha}{2} \quad (3)$$

Let us prove the following:

Claim 6.15. *There exists $u \in C_u(k-1)$, $d_v \in \{1, \dots, k-1\}$ and $v \in P \cap C_v(d_v)$ such that v is at distance at most k from u , and $|P \cap C_v(d_v)| = 2$.*

PROOF. Since P passes through C_u and C_v , let us consider a connected subgraph of P , which is a path denoted by P' , such that P' starts in C_u and ends in C_v . Let u be the last vertex in $P' \cap C_u(k-1)$ before P' passes through C_v . Since P' goes from $u \in C_u(k-1)$ to C_v , then it contains at least k other vertices after u (at least one in $C_v(d)$ for each $d \in \{0, \dots, k-1\}$). Let us denote by X the set of the k vertices following u in P' . All the vertices in X are at distance at most k from u . By definition of u , we have $X \cap (\bigcup_{0 \leq d \leq k-1} C_u(d)) = \emptyset$. And since $\alpha < \frac{2}{3}k$ and $|X| = k$, we have that $X \cap (\bigcup_{1 \leq d \leq k-1} C_v(d)) \neq \emptyset$. We may assume that for each vertex $x_1 \in X \cap \mathcal{B}_v$ there are at least two other distinct vertices $x_2, x_3 \in V(P) \cap \mathcal{B}_v$ such that x_1, x_2, x_3 are all the same distance from C_v .

Let D be the set of distances $d \in \{1, 2, \dots, k\}$ for which some $x \in X$ is at distance d from C_v (i.e. $x \in C_v(d)$). For each $d \in D$, we have $|P \cap C_v(d)| \geq 3$. And, by definition of X , D and α , we also have $|(\bigcup_{d \in D} C_v(d)) \cap X| \geq |X| - |X \setminus (\mathcal{B}_u \cup \mathcal{B}_v)| \geq k - \alpha$. Hence, $|(\bigcup_{d \in D} C_v(d)) \cap V(P)| \geq \max(3|D|, k - \alpha)$. Moreover for all $d \notin D$, we have $|V(P) \cap C_v(d)| \geq 2$. Thus, we get:

$$|V(P) \cap \mathcal{B}_v| = \left| \bigcup_{d \in D} C_v(d) \cap V(P) \right| + \left| \bigcup_{d \notin D} C_v(d) \cap V(P) \right| \geq \max(3|D|, k - \alpha) + 2(k - |D|) \quad (4)$$

There are two cases:

- if $|D| \geq \frac{k}{3} - \frac{\alpha}{2}$, then we have:

$$\begin{aligned} |\mathcal{B}_v \cap V(P)| &\geq 3|D| + 2(k - |D|) \\ &= 2k + |D| \\ &\geq \frac{7}{3}k - \frac{\alpha}{2} \end{aligned}$$

which contradicts (3).

- if $|D| \leq \frac{k}{3} - \frac{\alpha}{2}$, then we have:

$$\begin{aligned} |V(P) \cap \mathcal{B}_v| &\geq (k - \alpha) + 2(k - |D|) \\ &\geq (k - \alpha) + 2 \left(\frac{2k}{3} + \frac{\alpha}{2} \right) \\ &\geq \frac{7}{3}k \end{aligned}$$

which also contradicts (3).

In all cases, we get a contradiction. Thus, there exists $d_v \in \{1, \dots, k-1\}$ and a vertex $v \in X \cap C_v(d_v)$ such that $|P \cap C_v(d_v)| = 2$. ■

Using Claim 6.15, let us finish the proof by showing that u rejects at step (viii).

Let $u \in C_u(k-1)$, $d_v \in \{1, \dots, k-1\}$ and $v \in P \cap C_v(d_v)$ at distance at most k from u such that $|P \cap C_v(d_v)| = 2$. Let v' be the other vertex in $P \cap C_v(d_v)$. Then, u rejects at step (viii) of the verification. Indeed, there are several cases (see Figure 13 for an illustration):

- If P has both endpoints in $V \setminus (\bigcup_{d \leq d_v} C_v(d))$, the set of conditions (c) is satisfied.
- If P has both endpoints in $\bigcup_{d \leq d_v} C_v(d)$, the set of conditions (d) is satisfied.
- If P has exactly one endpoint in $\bigcup_{d \leq d_v} C_v(d)$, one of the sets of conditions (a) or (b) is satisfied. Indeed, Claim 6.15 proves more precisely that $u \in P \cap C_u(k-1)$, and that v is among the k next vertices after u in P . Let w be the vertex after v in P (which exists, otherwise v is an endpoint of P which would have both endpoints in $\bigcup_{d \leq d_v} C_v(d)$). If $w \in \bigcup_{d \leq d_v} C_v(d)$, then we are in case (a), else we are in case (b).

Conversely, let us show that if G is P_m -free, then no vertex rejects with the certificates described above. By definition, no vertex will reject in steps (i), (ii), (iii) and (vi). By contradiction, assume that there exists $u \in V$ which (iv), (v), (vii) or (viii).

- If u rejects at (iv), there exists an induced path P of length m in $G_{\leq u}$ such that all the vertices in $P \setminus V_{\leq u}$ are in distinct ECC_2 's. So P is also induced in G .
- If u rejects at step (v), with the same proof as in Theorem 1.10, we can reconstruct an induced path P of length m in G , which is a contradiction.
- If u rejects at step (vii), we can also reconstruct an induced path of length m in G . Indeed, let P_v (resp. P_w) be the path of length $\text{LongestPaths}[v]$ (resp. $\text{LongestPaths}[w]$) which starts at v (resp. w) and has all its other vertices in $\bigcup_{d < d_v} C_v(d)$ (resp. in $\bigcup_{d < d_w} C_w(d)$). Then, we obtain an induced path of length m by gluing P_v , P_{start} and P_w .
- Finally, if u rejects at step (viii), we can also construct an induced path P of length m in G , as depicted on Figure 13. In each case, the definition of $\text{LongestConstrainedPath}[v]$ ensures that the path obtained by gluing the parts is still induced in G .

In all cases, we obtain an induced path P of length m in G , which is a contradiction. Thus, if G is P_m -free, all vertices accept with the certificates described above. □

ACKNOWLEDGMENTS

L. Feuilloley thanks the participants of the PODC-DISC zulip chat for references about high-low degree techniques. This work was supported by ANR project GrR (ANR-18-CE40-0032). L. Cook was supported by the Institute for Basic Science (IBS-R029-C1).

REFERENCES

- [1] Nicolas Bousquet, Laurent Feuilloley, and Théo Pierron. 2021. Local Certification of Graph Decompositions and Applications to Minor-Free Classes. In *25th International Conference on Principles of Distributed Systems*, Vol. 217. 22:1–22:17. <https://doi.org/10.4230/LIPICS.OPODIS.2021.22>
- [2] Nicolas Bousquet, Laurent Feuilloley, and Sébastien Zeitoun. 2024. Local certification of local properties: tight bounds, trade-offs and new parameters. In *41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024)*, Vol. To appear.
- [3] Keren Censor-Hillel. 2022. Distributed Subgraph Finding: Progress and Challenges. *CoRR* abs/2203.06597 (2022). <https://doi.org/10.48550/ARXIV.2203.06597>
- [4] Keren Censor-Hillel, Orr Fischer, Tzvil Gonen, François Le Gall, Dean Leitersdorf, and Rotem Oshman. 2020. Fast Distributed Algorithms for Girth, Cycles and Small Subgraphs. In *34th International Symposium on Distributed Computing, DISC 2020 (LIPIcs, Vol. 179)*. 33:1–33:17. <https://doi.org/10.4230/LIPICS.DISC.2020.33>
- [5] Keren Censor-Hillel, Ami Paz, and Mor Perry. 2020. Approximate proof-labeling schemes. *Theor. Comput. Sci.* 811 (2020), 112–124. <https://doi.org/10.1016/J.TCS.2018.08.020>
- [6] Yi-Jun Chang, Seth Pettie, and Hengjie Zhang. 2019. Distributed Triangle Detection via Expander Decomposition. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*. 821–840. <https://doi.org/10.1137/1.9781611975482.51>
- [7] Yi-Jun Chang and Thatchaphol Saranurak. 2020. Deterministic Distributed Expander Decomposition and Routing with Applications in Distributed Derandomization. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, Sandy Irani (Ed.). 377–388. <https://doi.org/10.1109/FOCS46700.2020.00043>
- [8] Aristotelis Chaniotis, Linda Cook, Sepehr Hajebi, and Sophie Spirkl. 2024. Personal Communication.
- [9] Pierluigi Crescenzi, Pierre Fraigniaud, and Ami Paz. 2019. Trade-Offs in Distributed Interactive Proofs. In *33rd International Symposium on Distributed Computing, DISC 2019*, Jukka Suomela (Ed.), Vol. 146. 13:1–13:17. <https://doi.org/10.4230/LIPICS.DISC.2019.13>
- [10] Oscar Defrain, Louis Esperet, Aurélie Lagoutte, Pat Morin, and Jean-Florent Raymond. 2023. Local certification of geometric graph classes. *CoRR* abs/2311.16953 (2023). <https://doi.org/10.48550/ARXIV.2311.16953>
- [11] Andrew Drucker, Fabian Kuhn, and Rotem Oshman. 2014. On the power of the congested clique model. In *ACM Symposium on Principles of Distributed Computing, PODC '14*, Magnús M. Halldórsson and Shlomi Dolev (Eds.). 367–376. <https://doi.org/10.1145/2611462.2611493>
- [12] Talya Eden, Nimrod Fiat, Orr Fischer, Fabian Kuhn, and Rotem Oshman. 2022. Sublinear-time distributed algorithms for detecting small cliques and even cycles. *Distributed Comput.* 35, 3 (2022), 207–234. <https://doi.org/10.1007/S00446-021-00409-3>
- [13] Louis Esperet and Benjamin Lévêque. 2022. Local certification of graphs on surfaces. *Theor. Comput. Sci.* 909 (2022), 68–75. <https://doi.org/10.1016/J.TCS.2022.01.023>
- [14] Louis Esperet and Sergey Norin. 2022. Testability and Local Certification of Monotone Properties in Minor-Closed Classes. In *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022*, Vol. 229. 58:1–58:15. <https://doi.org/10.4230/LIPICS.ICALP.2022.58>
- [15] Guy Even, Orr Fischer, Pierre Fraigniaud, Tzvil Gonen, Reut Levi, Moti Medina, Pedro Montealegre, Dennis Olivetti, Rotem Oshman, Ivan Rapaport, and Ioan Todinca. 2017. Three Notes on Distributed Property Testing. In *31st International Symposium on Distributed Computing, DISC 2017*, Vol. 91. 15:1–15:30. <https://doi.org/10.4230/LIPICS.DISC.2017.15>
- [16] Laurent Feuilloley. 2021. Introduction to local certification. *Discret. Math. Theor. Comput. Sci.* 23, 3 (2021). <https://doi.org/10.46298/DMTCS.6280>
- [17] Laurent Feuilloley, Nicolas Bousquet, and Théo Pierron. 2022. What Can Be Certified Compactly? Compact local certification of MSO properties in tree-like graphs. In *PODC '22: ACM Symposium on Principles of Distributed Computing*. ACM, 131–140. <https://doi.org/10.1145/3519270.3538416>
- [18] Laurent Feuilloley, Pierre Fraigniaud, Juho Hirvonen, Ami Paz, and Mor Perry. 2021. Redundancy in distributed proofs. *Distributed Comput.* 34, 2 (2021), 113–132. <https://doi.org/10.1007/S00446-020-00386-Z>
- [19] Laurent Feuilloley, Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, Éric Rémila, and Ioan Todinca. 2021. Compact Distributed Certification of Planar Graphs. *Algorithmica* 83, 7 (2021), 2215–2244. <https://doi.org/10.1007/S00453-021-00823-W>
- [20] Laurent Feuilloley, Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, Éric Rémila, and Ioan Todinca. 2023. Local certification of graphs with bounded genus. *Discret. Appl. Math.* 325 (2023), 9–36. <https://doi.org/10.1016/J.DAM.2022.10.004>
- [21] Orr Fischer, Tzvil Gonen, Fabian Kuhn, and Rotem Oshman. 2018. Possibilities and Impossibilities for Distributed Subgraph Detection. In *Proceedings of the 30th Symposium on Parallelism in Algorithms and Architectures*, Christian Scheideler and Jeremy T. Fineman (Eds.). 153–162. <https://doi.org/10.1145/3210377.3210401>

- [22] Pierre Fraigniaud, Frédéric Mazoit, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. 2023. Distributed Certification for Classes of Dense Graphs. In *37th International Symposium on Distributed Computing, DISC 2023 (LIPIcs, Vol. 281)*. 20:1–20:17. <https://doi.org/10.4230/LIPICS.DISC.2023.20>
- [23] Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. 2022. A Meta-Theorem for Distributed Certification. In *Structural Information and Communication Complexity - 29th International Colloquium, SIROCCO 2022*, Merav Parter (Ed.), Vol. 13298. 116–134. https://doi.org/10.1007/978-3-031-09993-9_7
- [24] Mika Göös and Jukka Suomela. 2016. Locally Checkable Proofs in Distributed Computing. *Theory of Computing* 12, 19 (2016), 1–33. <https://doi.org/10.4086/toc.2016.v012a019>
- [25] Benjamin Jauregui, Pedro Montealegre, Diego Ramírez-Romero, and Ivan Rapaport. 2023. Local Certification of Some Geometric Intersection Graph Classes. *CoRR* abs/2309.04789 (2023). <https://doi.org/10.48550/ARXIV.2309.04789>
- [26] Janne H. Korhonen and Joel Rybicki. 2017. Deterministic Subgraph Detection in Broadcast CONGEST. In *21st International Conference on Principles of Distributed Systems, OPODIS 2017*, Vol. 95. 4:1–4:16. <https://doi.org/10.4230/LIPICS.OPODIS.2017.4>
- [27] Amos Korman, Shay Kutten, and David Peleg. 2010. Proof labeling schemes. *Distributed Comput.* 22, 4 (2010), 215–233. <https://doi.org/10.1007/S00446-010-0095-3>
- [28] Amir Nikabadi and Janne H. Korhonen. 2021. Beyond Distributed Subgraph Detection: Induced Subgraphs, Multicolored Problems and Graph Parameters. In *25th International Conference on Principles of Distributed Systems, OPODIS 2021*, Vol. 217. 15:1–15:18. <https://doi.org/10.4230/LIPICS.OPODIS.2021.15>
- [29] Rafail Ostrovsky, Mor Perry, and Will Rosenbaum. 2017. Space-Time Tradeoffs for Distributed Verification. In *Structural Information and Communication Complexity - 24th International Colloquium, SIROCCO 2017*, Vol. 10641. 53–70. https://doi.org/10.1007/978-3-319-72050-0_4
- [30] Boaz Patt-Shamir and Mor Perry. 2022. Proof-labeling schemes: Broadcast, unicast and in between. *Theor. Comput. Sci.* 923 (2022), 179–195. <https://doi.org/10.1016/J.TCS.2022.05.006>
- [31] Jukka Suomela. 2020. Landscape of Locality (Invited Talk). In *17th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2020 (LIPIcs, Vol. 162)*, Susanne Albers (Ed.). 2:1–2:1. <https://doi.org/10.4230/LIPICS.SWAT.2020.2>