



HAL
open science

Local certification of local properties: tight bounds, trade-offs and new parameters

Nicolas Bousquet, Laurent Feuilloley, Sébastien Zeitoun

► To cite this version:

Nicolas Bousquet, Laurent Feuilloley, Sébastien Zeitoun. Local certification of local properties: tight bounds, trade-offs and new parameters. 41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024), Mar 2024, Clermont - Ferrand, France. 10.4230/LIPIcs.STACS.2024.21 . hal-04440851v2

HAL Id: hal-04440851

<https://hal.science/hal-04440851v2>

Submitted on 11 Oct 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.



Distributed under a Creative Commons Attribution 4.0 International License


Local Certification of Local Properties: Tight Bounds, Trade-Offs and New Parameters

Nicolas Bousquet ✉ 

Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205 F-69622 Villeurbanne, France

Laurent Feuilloley ✉ 

Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, F-69622 Villeurbanne, France

Sébastien Zeitoun ✉ 

Univ Lyon, CNRS, INSA Lyon, UCBL, LIRIS, UMR5205, F-69622 Villeurbanne, France

Abstract

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. For many classic global properties, like checking the acyclicity of the network, the optimal size of the certificates depends on the size of the network, n . In this paper, we focus on properties for which the size of the certificates does not depend on n but on other parameters.

We focus on three such important properties and prove tight bounds for all of them. Namely, we prove that the optimal certification size is: $\Theta(\log k)$ for k -colorability (and even exactly $\lceil \log k \rceil$ bits in the anonymous model while previous works had only proved a 2-bit lower bound); $(1/2) \log t + o(\log t)$ for dominating sets at distance t (an unexpected and tighter-than-usual bound); and $\Theta(\log \Delta)$ for perfect matching in graphs of maximum degree Δ (the first non-trivial bound parameterized by Δ). We also prove some surprising upper bounds, for example, certifying the existence of a perfect matching in a planar graph can be done with only two bits. In addition, we explore various specific cases for these properties, in particular improving our understanding of the trade-off between locality of the verification and certificate size.

2012 ACM Subject Classification Theory of computation \rightarrow Distributed algorithms

Keywords and phrases Local certification, local properties, proof-labeling schemes, locally checkable proofs, optimal certification size, colorability, dominating set, perfect matching, fault-tolerance, graph structure

Digital Object Identifier 10.4230/LIPIcs.STACS.2024.21

Related Version *Full Version*: <https://arxiv.org/abs/2312.13702> [5]

Funding This work is supported by the ANR grant GrR (ANR-18-CE40-0032).

Acknowledgements The authors thank anonymous reviewers for useful comments.

1 Introduction

Local certification. Local certification is a topic at the intersection of locality and fault-tolerance in distributed computing. Very roughly, the main concern is to measure how much information the nodes of a network need to know in order to verify that the network satisfies a given property. Local certification originates from self-stabilization and is tightly related to the minimal memory needed to be sure that a self-stabilizing algorithm has reached a correct configuration locally. The topic is now studied independently, and the area has been very active during the last decade. We refer to the survey [14] for an introduction to the topic.

The standard model for local certification is the following. The nodes are first assigned labels, called *certificates*, and then every node looks at its certificate and the certificates of its neighbors, and decides to accept or reject. A certification scheme for a given property is correct if, for any network, the property is satisfied if and only if there exists a certificate



© Nicolas Bousquet, Laurent Feuilloley, and Sébastien Zeitoun;
licensed under Creative Commons License CC-BY 4.0

41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024).

Editors: Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov;

Article No. 21; pp. 21:1–21:18



Leibniz International Proceedings in Informatics

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



21:2 Local Certification of Local Properties

assignment such that all the nodes accept. (We discuss variations later, and give proper definitions in Section 3.) The usual measure of performance of a certification scheme is the maximum certificate size over all nodes, and all networks of a given size.

One of the fundamental results in local certification is that any property can be certified if the network is equipped with unique identifiers [21, 20], but this is at the expense of huge certificates. Indeed, the scheme consists in giving to every node the full map of the graph, which takes $\Theta(n^2)$ bits in n -node graphs. The question then is: when can we do better? There exist three typical certificate sizes. For some properties, *e.g.* related to graph isomorphism, $\Theta(n^2)$ bits is the best we can do [20]. For many natural properties, the optimal certificate size is $\Theta(\log n)$; for example most properties related to trees (acyclicity, spanning tree, BFS, and minimum spanning tree for small edge weights [21]). A recent research direction tries to capture precisely which properties have such *compact certification* (see [4, 19, 18]). Finally, some properties are *local* from the certification point of view, in the sense that the optimal certification size *does not depend on n* . This third type of property is the topic of this paper.

Local certification of local properties

Until recently, studying local properties has not been the focus of the community, since the usual parameter for measuring complexity is the network size. A recent paper by Ardévol Martínez, Caoduro, Feuilloley, Narboni, Pournajafi and Raymond [22] is the first to target this regime. We refer to [22] for the full list of motivations to study this topic, and we just highlight a few points here.

First, it will appear in this paper that the size of the certificates for local properties is often expressed as functions of parameters different from the number of nodes. Therefore, one should not (always) see these are constants. It has been highlighted before (see discussion before Open problem 4 in [14]) that we have basically no understanding of certification size expressed by other parameters than the number of nodes.

Second, *locally checkable languages* (LCL) are at the core of the study of the LOCAL and CONGEST models. These are basically the properties that are local from a certification point of view: the output can be checked by looking at all the balls of some constant radius. Certification is a way to question the encoding of LCLs. A typical example is coloring, for which one uses the colors as output of a construction algorithm, and as certificates for colorability certification. If there would exist a better certification, this would shed a new light on the encoding of this LCL. For example, the celebrated round elimination technique [25] is very sensitive to the problem encoding, and one could hope that feeding it with a different encoding could provide new bounds. In a more general perspective, we argue that just like understanding the complexity of LCLs, understanding certification of local properties is a fundamental topic.

In addition to these two general motivations, our work is guided by two open problems, the k -colorability question and the trade-off conjecture, that we detail now.

The k -colorability question

Let us first discuss the case of colorability, which will be central in this paper. The property we want to certify is that the graph is k -colorable, that is, one can assign colors from $\{1, \dots, k\}$ to vertices such that no two neighbors have the same color. It is straightforward to design a local certification with k certificates for this property: the certificates encode colors in $\{1, \dots, k\}$ and the nodes just have to check that there is no conflict.¹ This uses $O(\log k)$ bits. The natural open problem here is the following.

¹ A *conflict* being an edge whose two endpoints are colored the same.

► **Open problem 1** (Open problem 1 in [14]). *Is $\Theta(\log k)$ optimal for k -colorability certification?*

The first result on that question is the very recent paper [22] which establishes that one bit is not enough to certify k -colorability. This lower bound holds in the anonymous and in the proof-labeling scheme models, that we will define later. In a nutshell, the technique is an indistinguishability argument: assuming that there exists a 1-bit certification, one can argue about the number of 1s in a node neighborhood and take an accepting certification of some k -colorable graph to derive an accepting certification of a $k + 1$ -clique, which is a contradiction.

Interestingly, [22] also shows a case where the natural encoding is not the best certification. Namely, certifying a distance-2 3-coloring can be done with only 1 bit, while the obvious encoding of the colors takes 2 bits. Finally, let us mention that non- k -colorability, the complement property, is much harder to certify. Indeed, it is proved in [20] that one needs $\tilde{\Omega}(n^2)$ bits to certify non-3-colorability (where $\tilde{\Omega}$ hides inverse logarithmic factors).

Trade-off conjecture

We finish this general introduction, with yet another motivation to study local properties. The *trade-off conjecture*, first stated in [15], basically states that for any property, if the optimal certification size is s for the classic certification mechanism, where the nodes see their neighbors certificates, then it is in $O(s/d)$ if the vertices are allowed to see their whole neighborhood at distance d . This was proved to be true for many classic properties, and in many large graph classes [15, 17, 24]. Implicitly, the big-O of the conjecture refers to functions of n , but for local properties, the conjecture is interesting only if it refers to the parameters that appear in the certificate size.

► **Open problem 2** (Trade-off conjecture). *Consider a property with optimal certification size s at distance 1 (where s depends on the natural parameters of the problem). Is it true that if we allow the verification algorithm to look at distance d , then the optimal size is at most $\alpha \cdot s/d$ for some constant α ?*

The authors of [22] argue that the conjecture might actually be wrong for local properties. In other words, there might exist properties such that looking further in the graph is useless (unless you can see the whole graph), or at least not as useful as claimed in [15] (instead of being s/d the optimal size could be a less-decreasing function of d). We consider this question to be very intriguing and important to the study of locality, and we will discuss it several times in the paper.

A sample of local properties

Local properties have different behaviors, which prevented us to establish general theorems capturing all of them. Instead, we looked for a sample of properties widely studied in the distributed community having diverse behaviors, and such that many other properties would behave similarly to one of the sample. First, we chose to study the colorability question, for the reasons cited above. Second we looked at domination at distance t , where a set of nodes is selected, and we want to check that every node is at distance at most t from a selected node. This property has inputs and an external parameter, which makes it very different from colorability. Moreover, a dominating set distance t is a building block for many self-stabilizing algorithms. Third, we study the property of “having a perfect matching”. This differs from

21:4 Local Certification of Local Properties

the two first ones by being an edge-related problem instead of the node-related problem, and it appears that the key parameter there is the maximum degree, a new parameter for certification. One more motivation is that matchings are classic objects in distributed graph algorithms.

Organization of the paper

The paper is organized the following way. After this introduction, we give a detailed overview of the context, results. We also give the proof techniques of the main results. Then, after a definition section, there are technical sections that correspond to the local properties we study. For perfect matchings, the proofs of the results does not appear in this version due to the page limit. All the omitted proofs can be found in the full version [5]. The overview and the technical parts can be read independently. Readers interested in motivations, general discussions, proof ideas and comparison with previous techniques can read the first and cherry-pick specific proofs they are curious about in the second; while others will prefer to go directly to the model section and formal proofs. The overview and the technical part are organized in the same way to allow easy back-and-forth reading.

2 Overview of our results and techniques

Quick description of the models

In order to state the results, we need to informally define the different models of certification (see Section 3 for more formal definitions).

- In the anonymous model, the nodes have no identifiers/port-numbers.
- In the proof-labeling scheme model, every node has a unique identifier (encoded on $O(\log n)$ bits) but it cannot access the identifiers of the other nodes.
- In the locally checkable proof model, nodes have identifiers and can see the identifiers of the other nodes.

The anonymous model is the one for which we have the largest number of results. It is usually less considered in the literature, but argue that for local properties it is the most natural. See the discussion in Subsection 3.2.

The standard assumption is that the nodes can only see their neighbors. Since we are interested in the trade-off conjecture, we will also consider certification at distance d , where the view is the full neighborhood at distance d .

It is often handy to say that the certificates are given by a *prover*, that intuitively tries to convince the nodes that the network is correct (both on correct and incorrect instances).

Now that we are equipped with these notions, we will review our results and techniques, problem by problem.

2.1 Overview for colorability

In this subsection, we consider the *k-colorability* property, already mentioned, which states that the graph is *k*-colorable.

Two lower bounds for colorability

We have already discussed the colorability property, and cited Open problem 1. Our first result is in the anonymous model, where we strongly solve Open problem 1 by determining the *exact bound*: in the anonymous model it is necessary and sufficient to have k different certificates.

► **Theorem 3.** *For every $k \geq 2$, in the anonymous model where vertices can see at distance 1, k certificates are needed in order to certify that a graph is k -colorable. Therefore, in the anonymous model, exactly $\lceil \log(k) \rceil$ bits are needed to certify k -colorability.*

The upper bound is trivial since we can simply give the colors as certificates. The technique to establish this lower bound is a form of crossing technique. We take a large enough complete k -partite graph, thus a graph that is maximally k -colorable, in the sense that any edge added to it would make it non- k -colorable. The idea of the proof is to argue by counting, that for any certificate assignment that would make all nodes accept this graph, there must exist two edges that we can cross (that is replace $(a, b); (x, y)$ by $(a, x); (b, y)$ for example) such that all neighborhoods appearing in this instance did appear in the previous one, thus no node rejects. In addition, the graph obtained after this crossing is non- k -colorable, which is a contradiction with the correctness of the scheme.

Now in the most general model, we also give an (asymptotic) tight bound, fully answering Open question 1. We actually prove a more general lower bound parametrized by the verification distance.

► **Theorem 4.** *In the locally checkable proofs model, at least $\Omega(\log(k)/d)$ bits are needed to certify k -colorability when the vertices are allowed to see their neighborhoods at distance d .*

We describe the proof of this result in a communication complexity framework to provide more intuition, although the actual proof does not rely on any black-box result from communication complexity. The instances we use have the typical shape of communication complexity constructions: the graph has two parts that correspond to the players (left and right) and a part in the middle. The middle part has a large diameter (to be sure that left and right cannot communicate at distance d), and has two sets of k special vertices on the left (top left and bottom left), and two sets of k special vertices on the right (top right and bottom right). The part of the left (resp. right) player is an antimatching between the top left and bottom left (resp. top right and bottom right), where an antimatching is a complete bipartite graph in which a matching has been removed. The middle part has a very constrained structure whose role is to enforce that the graph is k -colorable, if and only if, the left and right antimatchings are intuitively mirrors one of the other. The idea is then that the information about the exact matchings on the left and right parts has to be transferred to some node, to be compared, and this can happen only via the certificates. There are $k!$ possible forms for the left and right antimatchings, (because there are $k!$ possible matching in a complete bipartite graph of size k basically). Thus, the information that has to be transferred from one side of the graph to the other has size $k \log k$ (via Stirling equivalent) and since our graph has cuts of order k , we get the $\Omega(\log k)$ lower bound.

Uniquely k -colorable graphs and other natural counterexample candidates

Our two lower bound constructions have in common to be very constrained, in a precise sense: for every vertex v , every ball centered in v of radius at least 2 admits a unique proper k -coloring (up to color renaming). In this case, we say that the graph is *uniquely k -colorable at distance d* (where d is the radius of the neighborhood). It is easy to see that if a graph is uniquely k -colorable at distance d , then either it has a unique k -coloring or it is not k -colorable. Intuitively, graphs with this property are hard for certification, since there is no slack in the coloring, thus the transfer of information between different parts of the graph cannot a priori be compressed. Perhaps surprisingly, we prove that for these graphs the trade-off conjecture does hold, even in the anonymous model.

► **Theorem 5.** *For every $d \leq \log k$, in the anonymous model where vertices can see at distance d , $O(\log k/d)$ bits are sufficient to certify that a uniquely k -colorable graph at distance $d - 2$ is k -colorable.*

Let us briefly explain the main ingredient of that proof. Since the coloring is locally unique, by looking far enough, a node can decide which other nodes are in the same color class in a k -coloring, if one exists globally. This is not enough to be sure that the graph is k -colorable, because the color classes might not coincide nicely. For example, every cycle is uniquely 2-colorable at distance 1, but this does not certify that the full cycle is 2-colorable. The key idea is that a node will recover the name of its color by gathering the bits of information spread on the nodes *of its own color class* at distance at most d from it. This allows to spread the information of the color classes on several vertices and then use smaller certificates. Slightly more formally, the certifier will assign to a well-chosen subset of nodes X a special certificate. The vertices of X are chosen far enough from each other so that we can store information of the different color classes on the nodes close to it. But they are also chosen not too far away from each other to be sure that all the vertices are close to a vertex of X . Now every node just have to perform the following verification: (i) if it is too far from any vertex of X , it rejects, (ii) it checks that its color is the same for all vertices of X close to it, (iii) it checks that for all the vertices of X close to it, the color associated to it is different from colors given to its neighbors.

After Theorem 5, a question is whether we can go down to a constant number of bits, or in other words, how large the constant of the big-O needs to be. We prove that at a distance $O(\log k)$, *one bit is sufficient* (in other words only two different certificates are needed). Note that it means in particular that we can avoid the special certificate of the proof of Theorem 5 by being very careful on how we represent each vertex of X and the colors classes around it (and proving that even if a node makes a mistake in the choice of X when some vertices are indistinguishable, its decision will be anyway correct).

► **Theorem 6.** *In the anonymous model where vertices can see to distance d , 2 certificates are enough to certify that a uniquely k -colorable graph at distance $d - 2$ is k -colorable, if $d \geq 9\lceil \log_2 k \rceil + 8$.*

The scheme for this is built on the same framework, but with more ideas and technicalities, and we refer to the technical part for the details.

Now that we proved that locally uniquely colorable graphs are not going to help, we wonder what would be a good candidate to disprove the trade-off conjecture (of course one might not exist, if the conjecture is true). We believe that for graphs that are *almost* locally uniquely colorable (that is, they have few correct colorings locally, up to color renaming) the proof of Theorem 5 could be adapted, with more layers of technicalities. Hence, one might go for graphs that have *many* possible colorings. This could make sense from a communication complexity point of view, because there also exist difficult problems with many correct pairs of inputs (*e.g.* the disjointness problem). What is problematic here is that unlike in communication complexity, in our case, on a k -colorable graph, the prover has the choice of the coloring, thus can choose one that is easier to encode compactly.

Since graphs that are very structured seem to admit a linear scaling, another approach could be to consider graphs with a more chaotic structure. Random graphs are natural candidates here, but even if we could come up with a satisfying definition for what it means to certify the colorability of a random graph, it is not clear that this would be hard. Indeed, there exist efficient algorithms to k -color k -colorable graphs (*e.g.* [27]) that exploit only the local structure of the graph, so one could even hope for a verification without certificates.

In the end, graphs with large girth and large chromatic number might be the right type of graphs to consider because they have both a large number of colorings locally and a rigid structure globally. Several constructions for these have been designed (see *e.g.* [13, 23, 10]), but they are randomized or complex, which makes their study rather challenging.

Certification versus output encoding

More generally, it is interesting to explore the links between (1) certifying that a given structure exists (*e.g.* a coloring here, but also a perfect matching a bit later in the paper) (2) explicitly describing it, and (3) implicitly describing it, that is giving enough information to recover the structure, but not directly (*e.g.* with fewer bits than the natural encoding). We refer to [3] for the related topic of distributed zero-knowledge proofs.

For colorability, the case of perfect graphs is a nice example of the discrepancy that can exist between these. A graph is *perfect* if its chromatic number is equal to the size of its maximum clique for every induced subgraph. Many classic graph classes are perfect, *e.g.* bipartite graphs, chordal graphs, comparability graphs (see [26] for a survey on perfect graphs). If we are promised to be in such a graph, the verification (at distance 2) is very easy: a vertex just has to check whether it belongs to a clique of size larger than k (in which case it rejects, otherwise it accepts). Thus, no certificates are needed for colorability certification, while it seems really difficult to have the nodes output a coloring without giving them quite a lot of information. Proving formally such a discrepancy would be a nice result.

Finally, let us mention yet another lower bound approach, related to problem encodings. One can note that if we have certification using s bits for a property, where the local verification algorithm runs in polynomial time in the network size, then we have a *centralized* decision algorithm for this property in time $s^n \text{poly}(n)$. Indeed, one can just enumerate all the possible certificate assignments of this size and check whether one is accepted. One could hope that having a too-good bound for certification would imply that some big conjecture (*e.g.* SETH, the strong exponential hypothesis) is wrong, and get a conditional lower bound.² Unfortunately, this does not help us much for certification, since there are known algorithms for computing the chromatic number of a graph in time $O(c^n)$ with $c < 3$ (*e.g.* in [7]).

2.2 Overview for domination

The *domination at distance t* property applies to graphs with inputs: every node should be labeled with 0 or 1, and every node should be at distance at most t from a node labeled 1. To avoid confusion, let us highlight that in the domination property, the inputs are part of the instance and are different from the certificates.

This problem is quite different from colorability, in several ways. First, it is a problem that is centered on inputs: for any graph there are inputs that are correct, so in some sense it is the inputs that are certified more than the graph itself. (This is actually closer to the original self-stabilizing motivation of certifying the output of an algorithm.) Second, the natural certification has a different flavor: we give every node its distance to the closest node of the dominating set (that is, a node labeled 1), and every node checks that these distances do make sense. One can then consider domination at distance t to be the local analogue of acyclicity, which is certified by providing the nodes of the network with the distance to the root.

² Such bounds are not common in certification, but not unseen, see [11].

Given this analogy with acyclicity, a $\Theta(\log t/d)$ optimal certificate size is expected, and indeed we prove it. But we go one step further by providing precise bounds on the number of different certificates. First, for the anonymous model at distance 1, we prove that the optimal number of different certificates is basically \sqrt{t} .

► **Theorem 7.** *Let $t \in \mathbb{N}^*$. In the anonymous model where vertices can see to distance 1, at least $\sqrt{t-1}$ different certificates are needed to certify a dominating set at distance t , even if the graphs considered are just paths and cycles.*

The lower bound is again quite expected: the proof for this kind of bound is based on arguing whether the same pair of certificates can appear several times on a path or not, thus the square root pops up naturally.

► **Theorem 8.** *In the anonymous model where vertices can see to distance 1, $3 \cdot \lceil \sqrt{t} \rceil$ certificates are sufficient to certify a dominating set at distance t .*

This upper bound is more surprising. It reveals that the natural encoding (consisting of giving the minimum distance to a labeled vertex) is not the best, and that the square root is not an artifact of the lower bound proof but is necessary. The proof of Theorem 8 is based on an elegant argument using *de Bruijn words*. Roughly, for some parameters k and n , a de Bruijn word is a (cyclic) word on an alphabet of size k , which contains all the factors (that is subwords of consecutive letters) of size n exactly once. The idea here is that instead of giving the certificate r to nodes at distance r from a node labeled 1, we will give them the r -th letter in a predefined de Bruijn word that corresponds to $n = 2$ and $k = \sqrt{t}$. Since every node will see its neighbors' certificates, thus a factor of size at least 2, it will be able to decode what is its position in the word, thus its distance to the 1-labeled node. The parametrization $k = \sqrt{t}$ ensures that the de Bruijn word of the correct length exists. Finally, we generalize these bounds to larger distances using generalizations of the techniques described above.

► **Theorem 9.** *In the anonymous model where vertices can see to distance $d < \frac{t}{2}$, at least $\frac{2^d \sqrt{t} - 2d + 1}{2}$ different certificates are needed to certify a dominating set at distance t , even if the graphs considered are just paths and cycles.*

► **Theorem 10.** *In the anonymous model where vertices are allowed to see to distance d , $O(d^{d+1} \sqrt{t})$ certificates are sufficient to certify a dominating set at distance t .*

2.3 Overview for perfect matchings

A graph has the *perfect matching property* if it has a perfect matching, that is, a set of edges such that every vertex belongs to exactly one such edge. It is yet another type of property, that differs from the others by the fact that it has no built-in parameter (no number k of colors, or distance t). As we will see, the relevant parameter here is the maximum degree of the graph.

Perfect matching certification upper bounds

The natural way to locally encode a matching in distributed computing is to make every node know which port number corresponds to a matched edge (if the vertex is matched). In the port-number model, this directly leads to a certification: give the relevant port number

to each node, and let them check the consistency. This takes $O(\log \Delta)$ bits per node, but it requires port numbers, and the ability for the nodes to know the port numbers of their neighbors.

Our first result is that we do not need port-numbers (nor any kind of initial symmetry-breaking). The strategy for the prover is the following. First, choose a perfect matching, and color the matched edges such that there are no two edges (u, v) and (w, z) of the same color with (v, w) being an edge of the graph. We call this a *matching coloring*. Then, give to every node the color of its matched edge. Then every node simply checks that it has exactly one neighbor with the same color.

► **Theorem 11.** *Let $k \in \mathbb{N}^*$. Let \mathcal{C} be a class of graphs such that, for every $G \in \mathcal{C}$, if G has a perfect matching then G has a k -matching coloring. Then, in the anonymous model at distance 1, k certificates are enough to certify the existence of a perfect matching in G for every $G \in \mathcal{C}$.*

We can easily show (see Lemma 28 in the full version [5]) that if G has a perfect matching then it admits a $(2\Delta - 1)$ -matching coloring.

► **Corollary 12.** *For every graph G , $2\Delta - 1$ certificates are enough to certify the existence of a perfect matching, in the anonymous model at distance 1.*

To better appreciate the lower bound of the next paragraph, let us mention a surprising result on the upper bound side. Note first that if we can get a matching coloring using fewer colors, then we automatically get a more compact certification. Now consider a planar graph with an arbitrary perfect matching. We claim that this perfect matching can be colored with only four colors, even though planar graphs can have an arbitrary large maximum degree. Indeed, if we contract the edges of the matching, we still have a planar graph, and we can color this graph with four colors, by the four color theorem. Now undoing the contraction, and giving to the matched edges the color of their contracted vertices, we get a proper matching coloring with only four colors. We also prove a more general result on minor-free and bounded treewidth graphs.

► **Corollary 13.** *In the anonymous model at distance 1:*

- *Only 2 bits are enough to certify the existence of a perfect matching for planar graphs.*
- *Only $O(\log k)$ bits are needed to certify the existence of a perfect matching in K_k -minor-free graphs.*
- *Only $\lceil \log_2(k + 1) \rceil$ bits are needed to certify the existence of a perfect matching in graphs of treewidth at most k .*

Note that, all our upper bounds follow from the existence of k -matching colorings. One can easily remark that, if instead of certifying a perfect matching one is simply interested in representing and certifying a matching, we can also do it with the same method by simply coloring all the unmatched vertices with an additional color. Since all the graph classes we mention in the upper bounds are closed under vertex deletion, our results ensure that the following holds: We can represent and certify matchings with 2Δ certificates for general graphs, 5 certificates for planar graphs and $t + 2$ certificates for graphs of treewidth at most t .

Perfect matching certification lower bounds

In the full version [5], we prove the following lower bound for perfect matching certification.

21:10 Local Certification of Local Properties

► **Theorem 14.** *For every $\Delta \geq 2$, in the anonymous model where vertices can see at distance 1, Δ different certificates are needed to certify the existence of a perfect matching for graphs of maximum degree Δ .*

Before we sketch the proof, let us note that there is no natural candidate for a lower bound: on the one hand, graphs that are sparse are ruled out by Corollary 13, and on the other hand many dense graph classes are known to always have perfect matchings, *e.g.* even cliques, or even random graphs with at least $n \log n$ edges [12].

The key to the proof is the notion of *half-graphs*. Half-graphs are bipartite graphs with vertices u_1, \dots, u_n and v_1, \dots, v_n such that u_i is linked to v_1, \dots, v_i . In such a graph, there exists a *unique* perfect matching, and it uses the edges (u_i, v_i) . Indeed, u_1 must be matched with v_1 , thus u_2 must be matched with v_2 etc. Now, we prove by counting argument that if such a graph is accepted with fewer than Δ certificates, we can create a graph without perfect matching that is also accepted. The idea is to take two copies of this certified graph, and then to carefully remove edges from within these graphs to add edges in between. Again, we refer to the technical section for the details.

Discussion of the parameter Δ

A remarkable feature of the perfect matching property is that the optimal certificate size is completely captured by the maximum degree (if we do not consider restricted graph classes). As far as we know, it is the first time that Δ appears as a natural parameter for local certification. This is interesting, since there are few results that use graph parameters other than the size of the graph to measure certificate size. To our knowledge, the only pure graph parameter that has been used before and that does not appear as a parameter of the problem is the girth, for approximate certification [11].³ It has been highlighted, *e.g.* in [14] (discussion before open problem 5), that developing a theory of parametrized certification is an interesting research direction.

Let us note however that it is maybe not very surprising that the maximum degree appears as a key parameter for matching-related problems, since celebrated papers have proved that it is central to the complexity of such problems in other models of computation, *e.g.* the LOCAL model [1, 6].

3 Model and definitions

3.1 Graphs

All the graphs we consider are finite, simple, and non-oriented. For completeness, let us recall the following classical graph definitions. Let $G = (V, E)$ be a graph, $u, v \in V$, $S \subseteq V$, $i \in \mathbb{N}$. The *distance between u and v* , denoted by $d(u, v)$, is the length (number of edges) of the shortest path from u to v . The *layer at distance i from u* , denoted by $N^i(u)$, is the set of vertices $v \in V$ such that $d(u, v) = i$. The *ball of radius i centered in u* , is $B(u, i) := \bigcup_{0 \leq j \leq i} N^j(u)$. The *closed (resp. open) neighborhood* of u is $N[u] := B(u, 1)$ (resp. $N^1(u)$). We can similarly define $d(u, S)$, $N(S)$ and $N[S]$ when S is a subset of vertices.

³ The maximum edge weight also appears for problems in weighted graphs, *e.g.* minimum spanning tree [21], max-weight matching [20, 8].

3.2 Certification

Let $G = (V, E)$ be a graph, and let C, I be non-empty sets. A *certificate function* of G (with certificates in C) is a mapping $c : V \rightarrow C$. An *identifier assignment* of G (with identifiers in I) is an injective mapping $Id : V \rightarrow I$.

► **Definition 15.** Let c be a certificate function of G and Id be an identifier assignment of G . Let $u \in V, d \in \mathbb{N}^*$. The view of u at distance d consists in all the information available at distance at most d from u , that is:

- the vertex u ;
- the graph with vertex set $B(u, d)$ and the edges $(v_1, v_2) \in E(G)$ such that $\{v_1, v_2\} \cap B(u, d-1) \neq \emptyset$;
- the restriction of c to $B(u, d)$;
- the restriction of Id to $B(u, d)$.

► **Remark 16.** For a vertex u , the subgraph induced by $B(u, d-1)$ is included in the view of u at distance d . However, the subgraph induced by $B(u, d)$ is *not* included in the view of u at distance d in general (because the view of u does not contain the edges between two vertices v_1 and v_2 which are both at distance exactly d from u).

A *verification algorithm (at distance d) in the locally checkable proof model* is a function which takes as input the view (at distance d) of a vertex, and outputs a decision, *accept* or *reject*. For a property \mathcal{P} on graphs, we say that there is a *certification for \mathcal{P}* using k certificates (resp. k bits) if C has size k (resp. 2^k), and if there exists a verification algorithm A such that for every graph G and every identifier assignment Id , G has property \mathcal{P} if and only if there exists a certificate function c such that A accepts for every $v \in V(G)$.

A *verification algorithm in the anonymous model* is defined in the exact same way that a verification algorithm in the locally checkable proof model, but vertices are not equipped with a unique identifier (or equivalently, the output is invariant with respect to the identifier assignment).

In this paper, we do not use the proof-labeling scheme model, where the node has access only its own ID, but it appears in a relevant previous work [22].

Discussion of the anonymous model

Note that the model we chose as our main model is the anonymous one. Indeed, all our results are for this model, except for the lower bound on colorability which we wanted to strengthen to the model with identifiers in order to fully solve Open Problem 1.

We think that the natural model for local properties is anonymous. Indeed, the main reason why identifiers are common in local certification is that they are often necessary, which is not the case for local properties. For example, certifying tree-like structures requires certifying that there is a unique connected component, and for this identifiers are needed. Note that in the LOCAL model, identifiers are also used to break symmetry, but in certification, the certificates can do this.

Moreover, the anonymous model is very common in the self-stabilizing literature (see *e.g.* [16, 2, 9]) which is the origin of local certification.

Third, in the paper, we draw a parallel between certification of local properties and LCLs, and the identifiers do not appear in the definition of LCLs.

Finally, in the cases known in the literature where anonymous and “with-ID” bounds match (e.g. acyclicity), the proofs are similar in spirit, except that the ID case involves more counting arguments (usually assuming that the ID interval is not of linear size), which implies loosing constants everywhere, getting results that are less crisp, and proofs that are more obfuscated.

4 Colorability certification

In this section, we will consider the certification of the k -colorability property. Let us recall that a graph G is said to be k -colorable if there exists *proper k -coloring* of G , that is, is a mapping $\varphi : V \rightarrow \{1, \dots, k\}$ such that for all $(u, v) \in E$, $\varphi(u) \neq \varphi(v)$.

For completeness, let us start by proving the following simple upper bound.

► **Proposition 17.** *In the anonymous model where vertices can see at distance 1, k -colorability can be certified with $\lceil \log k \rceil$ bits.*

Proof. For a graph $G = (V, E)$ which is k -colorable, the certificate function given by the prover is the following. The prover chooses a proper k -coloring φ of G , and assigns certificate $c(u) := \varphi(u)$ to every $u \in V$. The verification algorithm of every vertex u consists in checking if for every neighbor v , $c(u) \neq c(v)$. If it is the case, u accepts. Otherwise, u rejects. It is clear that the graph is accepted if and only if it is k -colorable. ◀

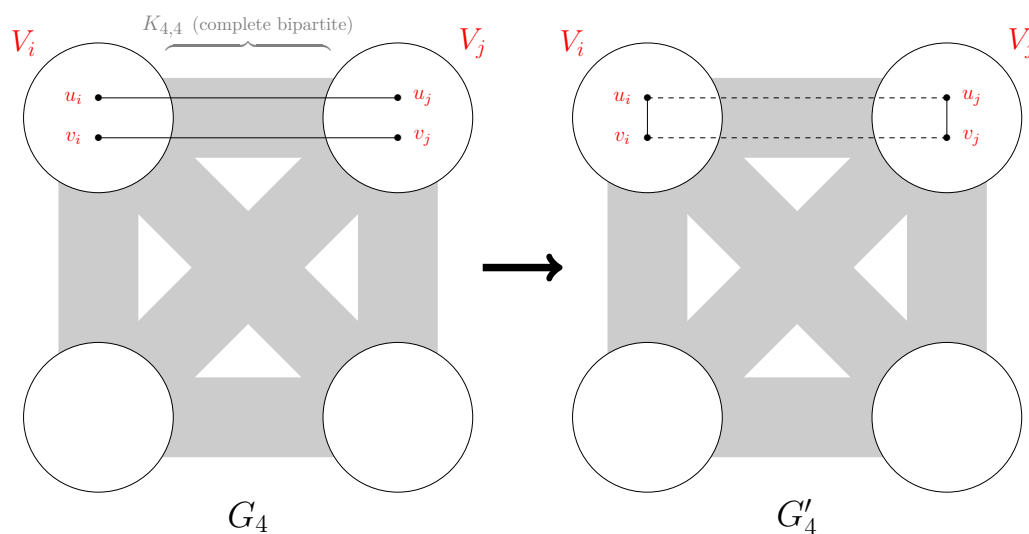
Lower bounds

► **Theorem 3.** *For every $k \geq 2$, in the anonymous model where vertices can see at distance 1, k certificates are needed in order to certify that a graph is k -colorable. Therefore, in the anonymous model, exactly $\lceil \log(k) \rceil$ bits are needed to certify k -colorability.*

Proof. Assume by contradiction that there exists a certification of k -colorability in the anonymous model using only $k - 1$ different certificates. The idea of the proof is to consider a specific k -colorable graph G_k , for which there must exist an accepting certification function. We will prove that we can flip (i.e. cross) two edges of G_k such that the resulting graph G'_k is not k -colorable and no vertex is able to detect it (meaning that the local view of each vertex will be unchanged). Thus, in the new graph G'_k , each vertex accepts, which is a contradiction since G'_k is not k -colorable.

Let us denote by G_k the complete k -partite graph, where each set has size $\max(k, 3)$. More formally, let V_1, \dots, V_k be k disjoint sets, each of size $\max(k, 3)$. Let G_k be the graph with vertex set $V = \bigcup_{i=1}^k V_i$ where (u, v) is an edge if and only if u and v do not belong to the same set V_i . Since $\{V_1, \dots, V_k\}$ is a partition of V into k independent sets, G_k is k -colorable. Thus, by hypothesis, there exists a certificate function $c : V \rightarrow \{1, \dots, k-1\}$ such that every vertex accepts. By the pigeonhole principle, for every $i \in \{1, \dots, k\}$ there exist two different vertices $u_i, v_i \in V_i$ such that $c(u_i) = c(v_i)$ (since each set has size at least k). Again, by the pigeonhole principle, there exist $i \neq j$ such that $c(u_i) = c(u_j)$. Thus, we get $c(u_i) = c(v_i) = c(u_j) = c(v_j)$, with $u_i, v_i \in V_i$ and $u_j, v_j \in V_j$. Let G'_k be the graph obtained from G_k by removing the two edges $(u_i, u_j), (v_i, v_j)$ and adding the two edges $(u_i, v_i), (u_j, v_j)$, as depicted on Figure 1.

We claim that, with the same certificate function c , all the vertices of G'_k accept. Indeed, the only vertices whose neighborhood have been modified between G_k and G'_k are u_i, v_i, u_j, v_j . So every vertex in $V \setminus \{u_i, v_i, u_j, v_j\}$ accepts. For u_i , the only difference between its



■ **Figure 1** Our constructions for the proof of Theorem 3, in the case of $k = 4$. The gray strips indicate complete bipartite graphs. The edges we are interested in appear explicitly.

neighborhood in G_k and G'_k is that u_j is replaced by v_i . Since $c(u_j) = c(v_i)$, the view of u_i is the same in G_k and G'_k , so u_i accepts in G'_k as well. Similarly, one can prove that v_i , u_j , and v_j accept in G'_k .

However, G'_k is not k -colorable. Indeed, assume by contradiction that it is, and let φ be a proper k -coloring of G'_k . For any $r \notin \{i, j\}$, let $w_r \in V_r$. Let w_j be a vertex of $V_j \setminus \{u_j, v_j\}$ which exists, since each set contains at least 3 vertices. Then, $K = \{w_1, \dots, w_{i-1}, w_{i+1}, \dots, w_k\}$ is a clique in G'_k , so the $(k - 1)$ vertices of K receive pairwise different colors in φ . Moreover, both u_i and v_i are complete to K . So if G'_k is k -colorable, then u_i and v_i have to be colored the same, which is a contradiction since $(u_i, v_i) \in E(G'_k)$. ◀

Obtaining lower bounds is usually harder in the locally checkable proofs model. In this more demanding model, we do not get the exact bound in terms of number certificates, but still we get a bound that would be considered optimal by the usual standards. Namely, we will prove that $\Omega(\log k)$ bits are needed to certify the k -colorability of a graph. More generally, we prove that, if vertices can see their neighborhoods at distance d , at least $\Omega(\log k/d)$ bits are needed to certify k -colorability. Note that the graph constructed in the proof of Theorem 3 has diameter two, thus no lower bound at distance at least 3 can be obtained from that construction. The lower bound of the following theorem is obtained with a completely different graph.


► **Theorem 4.** *In the locally checkable proofs model, at least $\Omega(\log(k)/d)$ bits are needed to certify k -colorability when the vertices are allowed to see their neighborhoods at distance d .*

Proof. For $r, s \geq 2$, let us denote by $P_{r,s}$ the graph on vertex set $V = \{1, \dots, r\} \times \{1, \dots, s\}$ with the following edges:

- for all $i \neq j \in \{1, \dots, r\}$ and all $p \in \{1, \dots, s\}$, $((i, p), (j, p)) \in E$
- for all $i \neq j \in \{1, \dots, r\}$ and all $p \in \{1, \dots, s - 1\}$, $((i, p), (j, p + 1)) \in E$


21:14 Local Certification of Local Properties

In other words, the graph $P_{r,s}$ is a sequence of s cliques $\mathcal{K}_1, \dots, \mathcal{K}_s$, each having size r , with an antimatching⁴ between \mathcal{K}_i and \mathcal{K}_{i+1} for all $i \leq s-1$. The second coordinate indicates the index of the clique the vertex belongs to, and for all (i, p) with $p \leq s-1$, (i, p) is the only vertex in \mathcal{K}_p which is not adjacent to $(i, p+1)$. For instance, the graph $P_{3,5}$ is depicted on Figure 2.



■ **Figure 2** The graph $P_{3,5}$. Here the cliques are triangles, and two consecutive cliques are linked by all possible edges except for the horizontal ones (the antimatching).

Now, for any pair of permutations σ, τ of $\{1, \dots, r\}$, let us denote by $G_{r,s}(\sigma, \tau)$ the graph obtained in the following way. We take two copies of $P_{r,s}$, and we denote their cliques by $\mathcal{K}_1, \dots, \mathcal{K}_s$ and $\mathcal{K}'_1, \dots, \mathcal{K}'_s$. We give $2rs$ different identifiers to the vertices of the two copies (the identifiers are fixed, they do not depend on σ and τ). Finally, we add the antimatching $\{(i, \sigma(i))\}_i$ between \mathcal{K}_1 and \mathcal{K}'_1 , and the antimatching $\{(i, \tau(i))\}_i$ between \mathcal{K}_s and \mathcal{K}'_s (see Figure 3 for an illustration).



■ **Figure 3** The graph $G_{3,5}(\sigma, \tau)$, where σ is the permutation $(1, 2)$ and τ the cycle $(1, 2, 3)$.

▷ **Claim 18.** The graph $G_{r,s}(\sigma, \tau)$ is r -colorable if and only if $\sigma = \tau$.

Proof. Since \mathcal{K}_1 is a clique of size r , in a proper coloring with r colors, every color appears exactly once inside \mathcal{K}_1 . Similarly, every color appears exactly once in a proper coloring of \mathcal{K}_2 . For all $i \in \{1, \dots, r\}$, the vertex $(i, 1)$ is a neighbor of every vertex in \mathcal{K}_2 except $(i, 2)$. Thus, the vertices $(i, 1)$ and $(i, 2)$ are colored the same. Hence, given a coloring of \mathcal{K}_1 , there exists a unique way to properly color \mathcal{K}_2 . This propagates along the cycle of cliques $\mathcal{K}_2, \dots, \mathcal{K}_s, \mathcal{K}'_s, \dots, \mathcal{K}'_1$, and it leads to a proper coloring of the whole graph if and only if the coloring of \mathcal{K}'_1 is compatible with the coloring of \mathcal{K}_1 , that is, if and only if $\sigma = \tau$. ◁

⁴ Remember that an antimatching between two sets of nodes is an edge set containing all the possible edges between the two sets, except for a matching.

Using Claim 18, we will deduce a lower bound on the number of bits needed to certify k -colorability, when vertices are allowed to see at distance d . We consider the construction above with parameters $r = k$ and $s = 2d$. Assume that m bits are sufficient. Then, for any permutation σ , there exists a certificate function $c_\sigma : V \rightarrow \{0, \dots, 2^m - 1\}$ such that all the vertices of $G_{k,2d}(\sigma, \sigma)$ accept. Suppose that there are two different permutations σ, τ such that the certification functions are the same: $c_\sigma = c_\tau$, that is, all vertices receive the same certificate in both instances. Then $G_{k,2d}(\sigma, \tau)$ would be accepted with this certificate function, since the view of every vertex would be identical as its view in $G_{k,2d}(\sigma, \sigma)$ (for vertices in the left part) or in $G_{k,2d}(\tau, \tau)$ (for vertices in the right part). This would be a contradiction because $G_{k,2d}(\sigma, \tau)$ is not k -colorable (by Claim 18).

Hence, all the certificate functions c_σ are different. In particular, there are no more permutations of $\{1, \dots, k\}$ than functions $V \rightarrow \{0, \dots, 2^m - 1\}$. The set V has size $4kd$ since it is made of two copies of a graph of size $k \times 2d$. Therefore, we get $k! \leq 2^{4mkd}$, leading to $m \geq \frac{\log_2(k!)}{4kd}$. Finally, since $\log_2(k!) = \Omega(k \log k)$, we get the result. \blacktriangleleft

► **Remark 19.** Two ingredients of the proof, the communication complexity insight and the antimatchings to propagate the coloring, have already been used in [20], to get a lower bound on the certification of not-3-colorable graphs.

5 Dominating sets at distance t

Most of the results concerning dominating sets at distance t can be found in the full version [5]. We simply give a proof of Theorem 8 with an elegant argument using de Bruijn words:

► **Theorem 8.** *In the anonymous model where vertices can see to distance 1, $3 \cdot \lceil \sqrt{t} \rceil$ certificates are sufficient to certify a dominating set at distance t .*

Before proving Theorem 8, we will define a few notions. Let A be an alphabet, and $\omega, \omega' \in A^*$. We say that ω is a *factor* of ω' if there exists a sequence of consecutive letters in ω' which is equal to ω . Let us now define *de Bruijn words*, whose existence is well-known.

► **Proposition 20.** *Let $k, n \in \mathbb{N}^*$, and A be an alphabet of size k . There exists a word $\omega \in A^*$ of length k^n , such that every word of A^n appears at most once as a factor of ω . Such a word ω is called a (k, n) -de Bruijn word.⁵*

Proof of Theorem 8. Let $\tau = \lceil \sqrt{t} \rceil$. Let us prove that 3τ certificates are sufficient to certify a dominating set at distance t (in the anonymous model, where vertices can see at distance 1). Let $A := \{1, \dots, \tau\}$. The certificates used in the scheme will be pairs in $C := \{0, 1, 2\} \times A$. For $(x, y) \in C$, let $\pi_1(x, y) := x$ and $\pi_2(x, y) := y$. Let $\omega' \in A^*$ be a $(\tau, 2)$ -de Bruijn word (which, by definition, has length at least t), and let us denote by $\omega = \omega_1 \dots \omega_t$ its prefix of length exactly t .

Let $G = (V, E)$ be a graph and $S \subseteq V$ be the set of labeled vertices. If S is dominating at distance t , the certificate function given by the prover is the following. The vertices of S are given an arbitrary certificate, and for every $u \in V \setminus S$ at distance i from S , the prover gives to u the certificate $c(u)$ which is $(i \bmod 3, \omega_i)$.

⁵ Usually, de Bruijn words are defined as words of length k^n such that each word of A^n appears at least once when ω is considered circularly. But due to the length of ω , each word of A^n actually appears exactly once circularly, so at most once if ω is not seen as a circular word.

21:16 Local Certification of Local Properties

The informal idea of the verification is the following one. In its certificate, every vertex u is given a letter of ω . By looking at its neighbors, u will be able to determine its position in ω (since a pair of letters defines a unique position in the de Bruijn word ω), which corresponds to its distance to S .

More formally, let c be a certificate function. Each vertex u checks the certificate as follows :

- (i) If $N[u] \cap S \neq \emptyset$, then u accepts.⁶
- (ii) Else, u checks that, for all $u', u'' \in N[u]$, if $\pi_1(c(u')) = \pi_1(c(u''))$ then $\pi_2(c(u')) = \pi_2(c(u''))$. If it is not the case, u rejects.
- (iii) Then, u checks if it has at least one neighbor v such that $\pi_1(c(v)) = \pi_1(c(u)) - 1 \pmod 3$, and if $\pi_2(c(v))\pi_2(c(u))$ is a factor of ω . If it is not the case, u rejects.
- (iv) Finally, for every neighbor w such that $\pi_1(c(w)) = \pi_1(c(u)) + 1 \pmod 3$, u checks if ω has $\pi_2(c(v))\pi_2(c(u))\pi_2(c(w))$ as a factor. If it is not the case, then u rejects.
- (v) If u did not reject at this point, it accepts.

It remains to show that there exists a certificate function such that all the vertices of G accept if and only if S is dominating at distance t . If S is a dominating set at distance t , then one can easily check that all the vertices accept with the certificates assigned by the prover as described previously. Note that with this certificate function, for step (ii), if $u', u'' \in N[u]$ satisfy $\pi_1(c(u')) = \pi_1(c(u''))$, then $d(u', S) = d(u'', S)$ so $\pi_2(c(u')) = \pi_2(c(u''))$.

For the converse, assume that G is accepted with some certificate function c . By (iii), every $u \in V$ such that no vertex is labeled in $N[u]$ should have a neighbor v such that $\pi_1(c(v)) = \pi_1(c(u)) - 1 \pmod 3$ and $\pi_2(c(v))\pi_2(c(u)) = \omega_\ell \omega_{\ell+1}$ for some ℓ . Note that since ω is a de Bruijn word, this ℓ is unique. Let us prove by induction on $\ell \in \{1, \dots, t-1\}$ that $d(u, S) \leq \ell + 1$.

- For $\ell = 1$, we have $\pi_2(c(v))\pi_2(c(u)) = \omega_1 \omega_2$. Let us prove that $d(u, S) \leq 2$. It is sufficient to prove that $d(v, S) \leq 1$. If v accepts at step (i), the conclusion holds. So we can assume that v accepts at step (v). By (iii), v has a neighbor v' such that $\pi_1(c(v')) = \pi_1(c(v)) - 1 \pmod 3$ and $\pi_2(c(v'))\pi_2(c(v))$ is a factor of ω . Since v does not reject at step (iv), $\pi_2(c(v'))\pi_2(c(v))\pi_2(c(u))$ is a factor of ω . So the letters $\omega_1 \omega_2 = \pi_2(c(v))\pi_2(c(u))$ appear at least twice as a factor of ω , which is a contradiction with the definition of de Bruijn words. Indeed, $\pi_2(c(v))\pi_2(c(u))$ are already the first two letters of ω . So if a factor $\pi_2(c(v'))\omega_1 \omega_2$ appears somewhere, $\omega_1 \omega_2$ must appear at least twice. Hence, v accepts at step (i), so $d(v, S) \leq 1$.
- Assume now that $\ell \geq 2$. To prove that $d(u, S) \leq \ell + 1$, it is sufficient to prove that $d(v, S) \leq \ell$. If v accepts at step (i) the conclusion holds. So we can assume that v accepts at step (v). By (iii), v has a neighbor v' such that $\pi_1(c(v')) = \pi_1(c(v)) - 1 \pmod 3$ and $\pi_2(c(v'))\pi_2(c(v))$ is a factor of ω . Since v does not reject at step (iv), $\pi_2(c(v'))\pi_2(c(v))\pi_2(c(u))$ is a factor of ω . Since the factor $\pi_2(c(v))\pi_2(c(u))$ appears exactly once in ω and $\pi_2(c(v))\pi_2(c(u)) = \omega_\ell \omega_{\ell+1}$, we have $\pi_2(c(v'))\pi_2(c(v)) = \omega_{\ell-1} \omega_\ell$ by (iv). By induction hypothesis, it implies that $d(v, S) \leq \ell$.

Thus, for any vertex $u \in V$, if u accepts at step (i) we have $d(u, S) \leq 1$, otherwise we have $d(u, S) \leq t$ by the previous induction. So S is indeed a dominating set at distance t . ◀

⁶ We could verify that $\pi_2(c(u))$ corresponds to a letter in the beginning of ω , but it is not necessary.

References

- 1 Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. *J. ACM*, 68(5):39:1–39:30, 2021. doi:10.1145/3461458.
- 2 Samuel Bernard, Stéphane Devismes, Maria Gradinariu Potop-Butucaru, and Sébastien Tixeuil. Optimal deterministic self-stabilizing vertex coloring in unidirectional anonymous networks. In *23rd IEEE International Symposium on Parallel and Distributed Processing, IPDPS 2009*, pages 1–8, 2009. doi:10.1109/IPDPS.2009.5161053.
- 3 Aviv Bick, Gillat Kol, and Rotem Oshman. Distributed zero-knowledge proofs over networks. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022*, pages 2426–2458. SIAM, 2022. doi:10.1137/1.9781611977073.97.
- 4 Nicolas Bousquet, Laurent Feuilloley, and Théo Pierron. 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*, pages 131–140, 2022. doi:10.1145/3519270.3538416.
- 5 Nicolas Bousquet, Laurent Feuilloley, and Sébastien Zeitoun. Local certification of local properties: tight bounds, trade-offs and new parameters, 2023. arXiv:2312.13702.
- 6 Sebastian Brandt and Dennis Olivetti. Truly tight-in- Δ bounds for bipartite maximal matching and variants. In *PODC '20: ACM Symposium on Principles of Distributed Computing*, pages 69–78, 2020. doi:10.1145/3382734.3405745.
- 7 Jesper Makhholm Byskov. Enumerating maximal independent sets with applications to graph colouring. *Oper. Res. Lett.*, 32(6):547–556, 2004. doi:10.1016/j.orl.2004.03.002.
- 8 Keren Censor-Hillel, Ami Paz, and Mor Perry. Approximate proof-labeling schemes. *Theor. Comput. Sci.*, 811:112–124, 2020. doi:10.1016/j.tcs.2018.08.020.
- 9 Johanne Cohen, Jonas Lefèvre, Khaled Maâmra, Laurence Pilard, and Devan Sohler. A self-stabilizing algorithm for maximal matching in anonymous networks. *Parallel Process. Lett.*, 26(4):1650016:1–1650016:17, 2016. doi:10.1142/S012962641650016X.
- 10 Amin Coja-Oghlan, Charilaos Efthymiou, and Samuel Hetterich. On the chromatic number of random regular graphs. *J. Comb. Theory, Ser. B*, 116:367–439, 2016. doi:10.1016/j.jctb.2015.09.006.
- 11 Yuval Emek and Yuval Gil. Twenty-two new approximate proof labeling schemes. In *34th International Symposium on Distributed Computing, DISC 2020*, volume 179, pages 20:1–20:14, 2020. doi:10.4230/LIPIcs.DISC.2020.20.
- 12 P Erdős and A Rényi. On the existence of a factor of degree one of a connected random graph. *Acta Mathematica Hungarica*, 17(3-4):359–368, 1966.
- 13 Paul Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.
- 14 Laurent Feuilloley. Introduction to local certification. *Discret. Math. Theor. Comput. Sci.*, 23(3), 2021. doi:10.46298/dmtcs.6280.
- 15 Laurent Feuilloley, Pierre Fraigniaud, Juho Hirvonen, Ami Paz, and Mor Perry. Redundancy in distributed proofs. *Distributed Comput.*, 34(2):113–132, 2021. doi:10.1007/s00446-020-00386-z.
- 16 Michael J. Fischer and Hong Jiang. Self-stabilizing leader election in networks of finite-state anonymous agents. In Alexander A. Shvartsman, editor, *Principles of Distributed Systems, 10th International Conference, OPODIS 2006*, volume 4305, pages 395–409, 2006. doi:10.1007/11945529_28.
- 17 Orr Fischer, Rotem Oshman, and Dana Shamir. Explicit space-time tradeoffs for proof labeling schemes in graphs with small separators. In *25th International Conference on Principles of Distributed Systems, OPODIS 2021*, volume 217 of *LIPIcs*, pages 21:1–21:22, 2021. doi:10.4230/LIPIcs.OPODIS.2021.21.
- 18 Pierre Fraigniaud, Frédéric Mazoit, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. Distributed certification for classes of dense graphs. *CoRR*, abs/2307.14292, 2023. doi:10.48550/arXiv.2307.14292.

- 19 Pierre Fraigniaud, Pedro Montealegre, Ivan Rapaport, and Ioan Todinca. A meta-theorem for distributed certification. In *Structural Information and Communication Complexity - 29th International Colloquium, SIROCCO 2022*, volume 13298, pages 116–134, 2022. doi:10.1007/978-3-031-09993-9_7.
- 20 Mika Göös and Jukka Suomela. Locally checkable proofs in distributed computing. *Theory Comput.*, 12(1):1–33, 2016. doi:10.4086/toc.2016.v012a019.
- 21 Amos Korman, Shay Kutten, and David Peleg. Proof labeling schemes. *Distributed Comput.*, 22(4):215–233, 2010. doi:10.1007/s00446-010-0095-3.
- 22 Virginia Ardévol Martínez, Marco Caoduro, Laurent Feuilloley, Jonathan Narboni, Pegah Pournajafi, and Jean-Florent Raymond. Lower bound for constant-size local certification. In *Stabilization, Safety, and Security of Distributed Systems - 24th International Symposium, SSS 2022*, volume 13751, pages 239–253, 2022. doi:10.1007/978-3-031-21017-4_16.
- 23 Moshe Morgenstern. Existence and explicit constructions of $q + 1$ regular ramanujan graphs for every prime power q . *J. Comb. Theory, Ser. B*, 62(1):44–62, 1994. doi:10.1006/jctb.1994.1054.
- 24 Rafail Ostrovsky, Mor Perry, and Will Rosenbaum. Space-time tradeoffs for distributed verification. In *Structural Information and Communication Complexity - 24th International Colloquium, SIROCCO 2017*, volume 10641, pages 53–70, 2017. doi:10.1007/978-3-319-72050-0_4.
- 25 Jukka Suomela. Using round elimination to understand locality. *SIGACT News*, 51(3):63–81, 2020. doi:10.1145/3427361.3427374.
- 26 Nicolas Trotignon. Perfect graphs: a survey. *CoRR*, abs/1301.5149, 2013. arXiv:1301.5149.
- 27 Jonathan S. Turner. Almost all k -colorable graphs are easy to color. *J. Algorithms*, 9(1):63–82, 1988. doi:10.1016/0196-6774(88)90005-3.