



HAL
open science

Renaming in distributed certification

Nicolas Bousquet, Louis Esperet, Laurent Feuilloley, Sébastien Zeitoun

► **To cite this version:**

Nicolas Bousquet, Louis Esperet, Laurent Feuilloley, Sébastien Zeitoun. Renaming in distributed certification. 2024. hal-04722289

HAL Id: hal-04722289

<https://hal.science/hal-04722289v1>

Preprint submitted on 4 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.

Renaming in distributed certification

Nicolas Bousquet

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

Louis Esperet

Laboratoire G-SCOP (CNRS, Univ. Grenoble Alpes), Grenoble, France

Laurent Feuilloley

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

Sébastien Zeitoun

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

Abstract

Local certification is the area of distributed network computing asking the following question: How to certify to the nodes of a network that a global property holds, if they are limited to a local verification?

In this area, it is often essential to have identifiers, that is, unique integers assigned to the nodes. In this short paper, we show how to reduce the range of the identifiers, in three different settings. More precisely, we show how to rename identifiers in the classical local certification setting, when we can (resp. cannot) choose the new identifiers, and we show how a global certificate can help to encode very compactly a new identifier assignment that is not injective in general, but still useful. We conclude with a number of applications of these three results.

2012 ACM Subject Classification Theory of computation → Distributed algorithms

Keywords and phrases Local Certification, identifiers, renaming, global certificate

Related Version *The introduction and the part about global certification heavily rely on a brief announcement that appeared at PODC 2024, by a subset of the authors: <https://dl.acm.org/doi/10.1145/3662158.3662781> [1]*

Acknowledgements We thank William Kuszmaul for fruitful discussion on hashing.

1 Introduction

The topic of distributed certification originates from self-stabilization in distributed computing, where the nodes of a network are provided with some pieces of information called *certificates*. These certificates can either be local (each node receives its own certificate), or global (there is a unique certificate, which is the same for all the nodes). The aim of the nodes is then to decide if the network satisfies a given property. To do so, each node should take its decision (accept or reject) based only on its local view in the network, which consists in its neighbors and their certificates. The correctness requirement for a certification scheme is the following one: for every network, the property is satisfied if, and only if, there exists an assignment of the certificates such that all the nodes accept. Unsurprisingly, the parameter we want to optimize is the size of the certificates, which is usually expressed as a function of n , the number of nodes in the network. For a given property \mathcal{P} , the optimal size of the certificates can be seen in some sense as a measure of the locality of \mathcal{P} : the smaller it is, the more local \mathcal{P} is. We refer to the survey [5] for an introduction to certification.

Renaming in local certification

In almost all the literature on distributed certification, one assumes that the nodes are equipped with unique identifiers, that is, every node is given a different integer, that it uses

as a name. The classical assumption in distributed graph algorithms is that the range of these identifiers is $[n^c] := \{1, \dots, n^c\}$ for some constant $c > 1$, where n is the number of nodes. This is equivalent to saying that identifiers are binary words of length $O(\log n)$. One reason one does not use integers from $[n]$, is that it gives too much power to the algorithm, for example electing a leader can be done with no communication by electing the vertex with identifier 1. The range $[n^c]$ is also the classical assumption in distributed certification, and having directly range $[n]$ allows to easily shave log factors from known results, as we will see at the end of the paper.

Therefore, in general, a natural question is whether we can perform a renaming procedure efficiently, that is, starting from identifiers in some arbitrary range $[M]$, and efficiently getting to range $[n]$ (which is equivalent to have an explicit bijection from the vertex set to $[n]$). In the context of certification, this translates to the question: suppose every node is given as input a new integer from the smaller range, can we certify that these integers form a proper identifier assignment? We prove two results related to this question.

► **Theorem 1.** *In any n -vertex graph, there exists an identifier renaming from range $[M]$ to range $[n]$ that can be locally certified with $O(\log M)$ bits.*

A direct consequence of this theorem is that for local certification with certificates of size $\Omega(\log n)$, assuming identifiers in $[n^c]$ or $[n]$ is equivalent. Indeed, when nodes are given identifiers in $[n^c]$, one can encode in the certificates a new identifier assignment from $[n]$ as well as its certification with $O(\log n)$ bits.

In Theorem 1, the new identifier assignment has a specific form that is dictated by the structure of the graph. In some cases we cannot choose this new identifier assignment, as it might be given by an adversary. We prove that such a renaming can be certified with significantly larger certificates.

► **Theorem 2.** *In any n -vertex graph, any identifier renaming from range $[M]$ to range $[n]$ can be locally certified with $O(n + \log M)$ bits.*

This theorem is a generalization of a similar result restricted to paths and identifier range $[n]$ that appeared in [7]. The same paper proves a lower bound of $\Omega(n)$. Note that the setting of Theorem 2 is equivalent to the PERMUTATION problem, where the nodes of the network are given integers from $[n]$ and they need to certify that these values form a permutation of $[n]$ (see [9] for more on this problem, in the setting of interactive certification protocols). We will see two simple applications of Theorems 1 and 2 in Section 4: for every ℓ , there are local certification schemes for K_ℓ -freeness (the property of not containing a clique of size ℓ) and diameter at most ℓ , with certificates of size $O(n)$. Both results are optimal, and improve earlier results by a logarithmic factor.

We will also show that Theorem 2 can be combined with hashing techniques described in the next paragraph to provide an efficient certification protocol for the DISTINCTNESS problem (see Section 4). In this problem, also studied in [9], the nodes are given values from $[m]$ and they have to certify that the values are pairwise distinct. PERMUTATION and DISTINCTNESS are useful in the context of isomorphism or non-isomorphism problems [9].

Global certification and perfect hashing

As mentioned above, there are two natural levels of locality in certification. In the first case, the certificates are local and the verification is local too; in the second case, the certificates are global, but the verification remains local. When speaking about local or global certification, we thus refer to the locality or globality of the certificates (and not of the verification, which

is always local). In general, these two levels of certification are related, because bounds for one can be derived from bounds for the other. Namely, a global certification scheme is a particular case of a local one, and conversely, a local certification scheme can be transformed into a global one by giving as global certificate the list of the local certificates of each node in the network (so that each node can simulate the local certification scheme by recovering its own local certificate from the global one, see [6] for more details). However, these generic transformations are usually not optimal.

The reason we mention the relation between global and local certification, is that identifiers play a significant role there. Consider the example of certifying that a graph is bipartite. This can be done in local certification by assigning to every node a single bit indicating its part in the bipartition. As observed in [6], we can translate this into a global certificate of size $\Theta(n \log n)$ consisting of all the pairs (identifiers, bit). Authors in [6] made the following conjecture (which is also discussed in [5], see Open Problem 9), in the standard case where the range of identifiers is polynomial in n :

► **Conjecture 3.** *The optimal size for global certification of BIPARTITENESS is $\Theta(n \log n)$.*

In other words, they conjectured that there is no way to compress the identifiers in the global certificate. Here, we prove that the conjecture is wrong: one can actually use only $\Theta(n)$ bits in this setting (in fact, we prove a more general result in terms of graph homomorphisms, see Theorem 12). To prove it, we used a technique of perfect hashing. Intuitively, for bipartiteness, the global certificate consists in (1) a bijection h between the set of identifiers and $[n]$, that can be encoded compactly, and (2) a list of n bits, where the i -th bit corresponds to the color in a proper two-coloring of the vertex whose identifier is mapped to i by the bijection h . The key ingredient is Theorem 7, which ensures that such a bijection h (called a *perfect hash function*) can be encoded compactly.

This technique of perfect hashing does not consist precisely in a renaming as we will define it in Section 2, but the idea is similar, namely compress the identifiers in range $[n]$. We also present another application of this hashing tool in Theorem 10, which gives an efficient local certification scheme for the DISTINCTNESS problem.

2 Models and definitions

For completeness, we now introduce some basic graph definitions. All the graphs we consider are finite, simple, non-oriented, and connected. Let G be a graph. The vertex set and edge set are denoted by $V(G)$ and $E(G)$, respectively. For every $u \in V(G)$, we denote by $N(u)$ the *open neighborhood* of u , which is set of vertices $v \in V$ such that $\{u, v\} \in E$. A *homomorphism* from a graph G to a graph H is a function $\varphi : V(G) \rightarrow V(H)$ such that, for every edge $\{u, v\} \in E(G)$, we have $\{\varphi(u), \varphi(v)\} \in E(H)$. Note that homomorphisms generalize colorings, since a graph G is k colorable if and only if exists a homomorphism from G to the clique of size k (in particular, G is bipartite if and only if there exists a homomorphism from G to an edge).

A *graph with inputs* is a pair (G, i) where G is a graph, and i is a mapping $V(G) \rightarrow \mathcal{I}$, where \mathcal{I} is a set called the set of *inputs* and i the *input function*. A *renaming with range \mathcal{I}* of a graph G is an *injective* input function $i : V(G) \rightarrow \mathcal{I}$. In particular, for any n -vertex graph G , a renaming with range $[n]$ is precisely a permutation $V(G) \rightarrow [n]$.

Now, let us define formally the model of certification. Let $M : \mathbb{N} \rightarrow \mathbb{N}$, called the *identifier range* (which is fixed: it is part of the framework for which certification schemes will be designed). Let $n = |V(G)|$. In the following, we just write M instead of $M(n)$ for the sake of

readability. An *identifier assignment* of G is an injective mapping $\text{Id} : V \rightarrow [M]$. Finally, let C be a set, called the set of *certificates*. We will define two different models of certification, namely *local* and *global* certification.

► **Definition 4** (Local certification). *In local certification, a certificate assignment for a graph G (possibly equipped with an input function i) is a mapping $c : V(G) \rightarrow C$. Given an identifier assignment Id , a certificate assignment c , and a vertex $u \in V(G)$, the view of u consists in all the information available in its neighborhood, that is:*

- its own identifier $\text{Id}(u)$;
- its own input $i(u)$ (if G is a graph with inputs);
- its own certificate $c(u)$;
- the set of identifiers, inputs (if G is a graph with inputs) and certificates of its neighbors, which is $\{(\text{Id}(v), i(v), c(v)) \mid v \in N(u)\}$.

► **Definition 5** (Global certification). *In global certification, a global certificate is a certificate $c \in C$. Given a graph G (possibly equipped with an input function i), an identifier assignment Id , a global certificate c , and a vertex $u \in V(G)$, the view of u consists in all the information available in its neighborhood, that is:*

- its own identifier $\text{Id}(u)$;
- its own input $i(u)$;
- the global certificate c ;
- the set of identifiers and inputs (if G is a graph with inputs) of its neighbors, which is $\{(\text{Id}(v), i(v)) \mid v \in N(u)\}$.

A *verification algorithm* is a function which takes as input the view of a vertex, and outputs a decision (*accept* or *reject*).

Let \mathcal{P} be a property on graphs (possibly with inputs). We say that there is a local certification scheme (resp. a global certification scheme) with size $s(n)$ and identifier range M if there exists a verification algorithm A such that, for all $n \in \mathbb{N}$, there exists set C of size $2^{s(n)}$ (equivalently, the elements of C can be seen as binary words of $s(n)$ bits) satisfying the following condition: for every graph G (possibly equipped with an input function i) with n vertices, G satisfies \mathcal{P} if and only if, for every identifier assignment Id with range M , there exists a certificate assignment $c : V(G) \rightarrow C$ (resp. a global certificate $c \in C$) such that A accepts on every vertex. We will often say that the certificates are given by a *prover* which, intuitively, try to convinces the vertices that \mathcal{P} is satisfied, but who can succeed only if G indeed satisfies \mathcal{P} .

Note that a verification algorithm is just a function, with no more requirements. In particular, it does not have to be decidable. However, in practice, when designing a certification scheme to prove upper bounds, it turns out to be decidable and often computable in polynomial time. The fact that no assumptions are made on this verification function in the definition just strengthens the results when proving lower bounds, by showing that it does not come from computational limits.

We conclude these preliminaries with the definition of perfect hashing and a classical result on the size of perfect hash families.

► **Definition 6.** *Let $k, \ell \in \mathbb{N}$ with $k \leq \ell$, and let H be a set of functions $[\ell] \rightarrow [k]$.*

1. *A function $h \in H$ is a perfect hash function for $S \subseteq [\ell]$ if $h(x) \neq h(y)$ for all $x, y \in S$, $x \neq y$.*

2. The family of functions H is a (k, ℓ) -perfect hash family if, for every $S \subseteq [\ell]$ with $|S| = k$, there exists $h \in H$ which is perfect for S .

We will need the following Theorem 7 (see e.g. [8] for a proof).

► **Theorem 7.** *Let $k, \ell \in \mathbb{N}$ with $k \leq \ell$. There exists a (k, ℓ) -perfect hash family $H_{k, \ell}$ which has size $\lceil ke^k \log \ell \rceil$.*

A consequence of Theorem 7 is that, to encode a hash function $h : [M] \rightarrow [n]$ which is perfect for a given set $S \subseteq M$ of size n (for instance, the set of identifiers of a n -vertex graph), we just need $\log \lceil ne^n \log M \rceil = O(n + \log \log M)$ bits. See Theorems 10 and 12 for examples of applications.

3 Renaming theorems

3.1 Certifying some renaming to $[n]$ with $O(\log M)$ bits

► **Theorem 1.** *In any n -vertex graph, there exists an identifier renaming from range $[M]$ to range $[n]$ that can be locally certified with $O(\log M)$ bits.*

Proof. More precisely, we will prove that the following holds: there exists a local certification scheme with $O(\log M)$ bits such that

1. if all the vertices accept with some certificate assignment, then the renaming written in the input is a correct renaming of range $[n]$, and
2. there exists a renaming of range $[n]$ such that all the vertices accept with some certificate assignment.


Recall that for an n -vertex graph G , a renaming with range $[n]$ is precisely a permutation $V(G) \rightarrow [n]$. So the two items above can be rewritten as: (1) if all vertices accept, then the input is a permutation of $[n]$ and (2) there exists a permutation of $[n]$ such that if it is given in input, then all vertices accept.

Description of a renaming that can be certified with $O(\log M)$ bits.

Let G be an n -vertex graph. For every $u \in V(G)$, let us denote by Id_u the original identifier of u , and by $\text{Renaming}(u)$ the renaming of u that is written in its input. Let us describe a renaming in range $[n]$ that can be certified using $O(\log M)$ bits. Let \mathcal{T} be a spanning tree of G , rooted at some vertex r . For every vertex $u \in V(G)$, let us denote by \mathcal{T}_u the subtree of \mathcal{T} rooted at u , and by $|\mathcal{T}_u|$ its number of vertices. We set $\text{Renaming}(r) := 1$. Now, let $u \in V(G)$ with children v_1, \dots, v_k in \mathcal{T} ordered by increasing (original) identifiers $\text{Id}_{v_1}, \dots, \text{Id}_{v_k}$. We define $\text{Renaming}(v_1) := \text{Renaming}(u) + 1$, and for every $i \in \{2, \dots, k\}$, $\text{Renaming}(v_i) := \text{Renaming}(v_{i-1}) + |\mathcal{T}_{v_{i-1}}|$. In other words, this renaming corresponds to the order in which the vertices would be visited in a depth-first-search in \mathcal{T} , starting at r . See Figure 1 for an example.

Certification of the renaming.

Let us describe the certificates given by the prover to the vertices. Let \mathcal{T} be the spanning tree of G which gives the renaming described previously. For every vertex $u \in V(G)$, we denote by \mathcal{T}_u the subtree of \mathcal{T} rooted at u . The certificate of every vertex u consists in four parts, denoted by $\text{Root}[u]$, $\text{Parent}[u]$, $\text{Distance}[u]$, $\text{Subtree}[u]$, which are defined as follows.



■ **Figure 1** Illustration of the renaming. The edges which are drawn here are those of \mathcal{T} . For each internal node, its children are written in increasing order of the original identifiers. The integer in every node u corresponds to $\text{Renaming}(u)$. The red integer near a node u corresponds to $|\mathcal{T}_u|$. (The original identifiers are omitted.)

- In $\text{Root}[u]$, the prover writes the identifier Id_r of r .
- In $\text{Parent}[u]$, the prover writes the identifier Id_w of the parent w of u in \mathcal{T} (if $u \neq r$) and Id_r otherwise.
- In $\text{Distance}[u]$, the prover writes the distance from u to r in \mathcal{T} .
- In $\text{Subtree}[u]$, the prover writes the number of vertices in \mathcal{T}_u .

Note that all of these four parts can be encoded with $O(\log M)$ bits (because $n \leq M$). Thus, the overall size of the certificates is $O(\log M)$.

Verification of the certificates.

Let us now explain how the vertices check the correctness of their certificates. If u and v are two neighbors such that $\text{Parent}[v] = \text{Id}_u$, we say that u is a *parent* of v , or equivalently that v is a *child* of u . If u does not have any children, we say that u is a *leaf*. The verification procedure of every vertex $u \in V(G)$ consists in several steps, which are the following ones. At each step, if the verification fails, u rejects. If u did not reject at any step, it finally accepts.

- (i) First, u checks that all its neighbors have the same root, that is $\text{Root}[u] = \text{Root}[v]$ for every neighbor v of u .
If no vertex rejects at this point, then Root is the same in the certificates of all the vertices, so in the following we will denote by Root this common value written in the certificates.
- (ii) If $\text{Id}_u \neq \text{Root}$, u checks that there is a vertex w in its neighborhood which is its parent, and it also checks that $\text{Distance}[w] = \text{Distance}[u] - 1$.
- (iii) If u is a leaf, it checks that $\text{Subtree}[u] = 1$. Otherwise, u checks that $\text{Subtree}[u] = 1 + \sum_{i=1}^k \text{Subtree}[v_i]$, where v_1, \dots, v_k are the children of u .
- (iv) Finally, u checks the correctness of the renaming. Namely, if $\text{Id}_u = \text{Root}$, u checks that $\text{Renaming}[u] = 1$. If u is not a leaf and has v_1, \dots, v_k as children, where the identifiers $\text{Id}_{v_1}, \dots, \text{Id}_{v_k}$ are in increasing order, it checks that $\text{Renaming}[v_1] = \text{Renaming}[u] + 1$, and that for all $i \in \{2, \dots, k\}$, $\text{Renaming}[v_i] = \text{Renaming}[v_{i-1}] + \text{Subtree}[v_{i-1}]$.

Correctness.

Let us prove the correctness of this certification scheme. First, assume that the renaming is the one described above, and that the prover gives the correct certificates to the vertices. Then, by definition of the renaming and the certificates, no vertex will reject in the verification procedure, so all the vertices accept.

Conversely, assume that all the vertices accept with some renaming and some certificates, and let us check that the renaming is correct. First, note that G contains a vertex having the identifier `Root`. Let r be the vertex of G such that $\text{Distance}[r]$ is minimal in G . Since r accepts, we have $\text{Id}_r = \text{Root}$, else r would reject at step (ii). Then, we can reconstruct the spanning tree \mathcal{T} that the prover used to assign the certificates: the root is r , and a vertex v is a child of u if u and v are neighbors and $\text{Parent}[v] = \text{Id}_u$. Step (ii) ensures that this definition leads to a correct spanning tree \mathcal{T} . Then, step (iii) ensures that for each vertex u , we have $\text{Subtree}[u] = |\mathcal{T}_u|$. Finally, step (iv) guarantees that `Renaming` corresponds the order in which the vertices would be visited in a depth-first search starting at r , which is indeed a correct renaming in $[n]$. ◀

3.2 Certifying any renaming to $[n]$ with $O(n + \log M)$ bits

► **Theorem 2.** *In any n -vertex graph, any identifier renaming from range $[M]$ to range $[n]$ can be locally certified with $O(n + \log M)$ bits.*

Proof. More precisely, what we prove is the following: there exists a local certification scheme with $O(n + \log M)$ bits such that

1. if all the vertices accept with some certificate assignment, then the renaming written in the input is a correct renaming of range $[n]$, and
2. for all renamings of range $[n]$, all the vertices accept with some certificate assignment.

As before, the two items above can be rewritten in the language of permutation as follows: (1) if all vertices accept, then the input is a permutation of $[n]$ and (2) for any permutation of $[n]$ in input, all vertices accept.

As in the proof of Theorem 1, for every $u \in V(G)$, we denote by $\text{Renaming}[u]$ the renaming of u which is written in its input.

Certification of the renaming.

Let us describe the certificates given by the prover to the vertices. The prover chooses a spanning tree \mathcal{T} of G rooted at some vertex r , and the certificate of every vertex u consists in four parts $\text{Root}[u]$, $\text{Parent}[u]$, $\text{Distance}[u]$, $\text{Subtree}[u]$. The parts $\text{Root}[u]$, $\text{Parent}[u]$ and $\text{Distance}[u]$ are defined exactly as in the proof of Theorem 1, and only $\text{Subtree}[u]$ differs. In $\text{Subtree}[u]$, the prover writes a binary vector of n bits whose j -th bit is equal to 1 if and only if there exists a vertex $v \in \mathcal{T}_u$ such that $\text{Renaming}[v] = j$.

The three first parts of the certificates have size $O(\log M)$, and the `Subtree` part has size $O(n)$, so the overall size of the certificates is $O(\log M + n)$.

Verification of the certificates.

Let us now explain how the vertices check the correctness of their certificates. The definitions of *parent*, *child* and *leaf* are the same as in the proof of Theorem 1. As in the proof of Theorem 1, the verification procedure consists in several steps. Steps (i) and (ii) are identical, and there are finally three last steps (iii), (iv) and (v) which are the following:

- (iii) u checks that the **Subtree** vectors of its neighbors have the same size, that is, for every neighbor v , $|\text{Subtree}[u]| = |\text{Subtree}[v]|$.
- (iv) If u is a leaf, it checks that $\text{Subtree}[u]_j = 1$ if and only if $j = \text{Renaming}[u]$ for every integer j . Otherwise, let us denote by v_1, \dots, v_k the children of u . For every integer j , the vertex u checks that one of the three following cases holds:
 - $\text{Subtree}[u]_j = 0$, $\text{Renaming}[u] \neq j$, and for all $i \in [k]$, $\text{Subtree}[v_i]_j = 0$, or
 - $\text{Subtree}[u]_j = 1$, $\text{Renaming}[u] \neq j$, and there exists a unique $i \in [k]$ such that $\text{Subtree}[v_i]_j = 1$, or
 - $\text{Subtree}[u]_j = 1$, $\text{Renaming}[u] = j$, and for all $i \in [k]$, $\text{Subtree}[v_i]_j = 0$.
- (v) If $\text{Id}_u = \text{Root}$, u checks that $\text{Subtree}[u]_j = 1$ for all index j .

Correctness.

Let us prove the correctness of this certification scheme. First, as in the proof of Theorem 1, if the renaming is correct and if the prover gives the certificates as described above, no vertex will reject by definition.

Conversely, assume that all the vertices accept with some renaming and some certificates, and let us check that the renaming is a correct one. Exactly as in the proof of Theorem 1, step (ii) ensures that there exists a vertex r such that $\text{Id}_r = \text{Root}$, and we can reconstruct a spanning tree \mathcal{T} rooted at r that the prover used to assign the certificates. Then, step (iii) ensures that **Subtree** has the same length in the certificates of all the vertices. Thanks to step (iv), we can note that the number of bits equals to 1 in **Subtree** $[u]$ is equal to the number of vertices in \mathcal{T}_u , and since all the bits of **Subtree** $[r]$ are equal to 1 because of step (v), then the length of **Subtree** is equal to n . Step (iv) also ensures $\text{Subtree}[u]_j$ is equal to 1 if and only if there exists $v \in \mathcal{T}_u$ such that $\text{Renaming}[v] = r$. Since all the bits of **Subtree** $[r]$ are equal to 1, it implies that for every $j \in \{1, \dots, n\}$, there exists a vertex u such that $\text{Renaming}[u] = j$. By cardinality, there is only one such vertex u . Thus, it is a correct renaming in the range $[n]$. ◀

4 Applications

4.1 Obtaining optimality in the $\Theta(n)$ -bit regime

In this subsection we focus on local certification and prove that our results can be used to obtain optimal certification schemes for several classical problems.

K_ℓ -freeness

It was proved¹ in [3, Theorem 15] that for any fixed $\ell \geq 4$, locally certifying that an n -vertex graph G is K_ℓ -free (meaning that G does not contain K_ℓ as a subgraph) requires certificates of size $\Omega(n)$. On the other hand, K_ℓ -freeness can easily be locally certified with certificates of $O(n \log n)$ bits: it suffices to assign to each vertex the list of (identifiers) of its neighbors as a certificate, and then the vertices can readily check the validity of these certificates, and the local view of each vertex v is enough to compute the size of a maximum clique containing v in G (and thus certifying that G is K_ℓ -free, if it is the case). Storing the list of identifiers of its neighbors can cost up to $\Omega(n \log n)$ bits in the worst case, so with this technique it does

¹ The model considered in [3] is a bit different from local certification, but their proof can easily be adapted to work in this model.

not seem to be possible to match the lower bound of $\Omega(n)$ from [3]. We now explain how to obtain an optimal local certification scheme for K_ℓ -freeness (ℓ fixed) with certificates of size $O(n)$, as a simple consequence of Theorem 1 or 2.

► **Theorem 8.** *For any fixed ℓ , K_ℓ -freeness can be locally certified with certificates of $O(n)$ bits in n -vertex graphs (and this is optimal).*

Proof. Start by renaming the identifiers of G in the range $[n]$ using Theorem 1 or 2. In either case, the renaming can be locally certified with $O(n)$ bits. Now, the certificate of u consists of a binary vector of size $[n]$ whose j th coordinate is 1 if and only if u is adjacent to the vertex whose new identifier is j . That is, we encode the neighborhood of u with a vector of size n . As explained above, this information is enough to decide whether G is K_ℓ -free, as desired.

The fact that the size of the certificates is optimal was proved in [3]. ◀

Diameter $\leq \ell$

It is proved in [2] that, for any constant ℓ , having diameter at most ℓ can be certified locally with certificates of $O(n \log n)$ bits. Moreover, they obtained a lower bound of $\Omega(n)$ bits on the certificates for this problem. We obtain the following improvement.

► **Theorem 9.** *For any fixed ℓ , having diameter at most ℓ can be locally certified with certificates of $O(n)$ bits in n -vertex graphs (and this is optimal).*

Proof. The upper bound technique in [2] is the following. In a correct instance, every node v is given as a certificate a table of size n . Each cell corresponds to a every node w of the graph and contains the identifier of w and the distance from v to w . Checking these tables consists in checking in parallel n BFS trees. More formally, for every w , v checks that at least one of its neighbors has distance one less to w , none have distance two or more less and, if $d(v, w)$ is at least two, v checks that (v, w) is not an edge. Every cell uses $O(\log n + \log \ell) = O(\log n)$ bits, hence the result.

In our scheme, we start by renaming the identifiers of G in the range $[n]$ using Theorem 1 or 2. In either case, the renaming can be locally certified with $O(n)$ bits. Then on a correct instance, the prover assigns to every node v a table T_v of n cells, where the cell $T_v[i]$ contains the distance from v to the node with new identifier i . The verification can be done as in [2], and the size of the certificates is $O(n)$. ◀

Distinctness

In the introduction we have mentioned the DISTINCTNESS problem: given a graph G with identifiers of range $[M]$, and an input function i of range $[m]$, the goal is to certify that the inputs of all the vertices are distinct (i.e. that i is injective). See [9] for a study of this problem in the interactive version of local certification. We prove the following result:

► **Theorem 10.** *The DISTINCTNESS property can be locally certified in n -vertex graphs with identifier range $[M]$ and input range $[m]$ with certificates of size $O(n + \log M + \log \log m)$.*

Note that Theorem 10 is in fact a generalization of Theorem 2. Indeed, Theorem 2 is just the particular case of Theorem 10 where $m = n$.

Proof. The certificates given by the prover to the vertices consist in two parts. First, the prover chooses a hash function $h : [m] \rightarrow [n]$ which is perfect for the set of inputs

$\{i(u) \mid u \in V(G)\}$, and writes this hash function in the certificate of every vertex. By Theorem 7, it uses $O(n + \log \log m)$ bits. Then, the prover uses the certification of Theorem 2 to certify that $h \circ i$ is a correct renaming, which uses $O(n + \log M)$ bits. In total, the certificates are of size $O(n + \log M + \log \log m)$. The verification procedure of the vertices just consists in checking if they all receive the same hash function h , and to perform the verification of Theorem 2 to check that $h \circ i$ is indeed injective. \blacktriangleleft

4.2 Global certification of homomorphism

In this subsection, we focus on global certification. The property we want to certify is H -HOMOMORPHISM (the existence of a homomorphism to a given graph H). A particular case which has already been studied in [6] is BIPARTITENESS (it corresponds to the case where H is the clique on two vertices). Note that there exists a local certification scheme for BIPARTITENESS using only one bit per vertex (where the certificate is the color in a proper two-coloring, and the verification of every node just consists in checking that it received a different color from all its neighbors). However, with a global certificate, it is less clear how to certify it optimally. In [6], the authors proved the following upper and lower bounds:

► **Theorem 11.** *If s denotes the optimal certificate size for the global certification of BIPARTITENESS for n -vertex graphs whose identifiers are in the range $[M]$, then:*

$$s = \Omega(n + \log \log M) \quad \text{and} \quad s = O(\min\{M, n \log M\})$$

The authors of [6] also made the conjecture that their lower bound can be improved to match their upper bound. Namely, in the standard case where $M = O(n^c)$, they conjectured the following:

► **Conjecture 3.** *The optimal size for global certification of BIPARTITENESS is $\Theta(n \log n)$.*

Here, we disprove this conjecture, by proving the following stronger result. The key ingredient used in the proof is perfect hashing (Theorem 7).

► **Theorem 12.** *For any graph H , there exists a global certification scheme for H -HOMOMORPHISM for n -vertex graphs whose identifiers are in the range $[M]$, with a certificate of size $O(n \log |V(H)| + \log \log M)$.*

► **Corollary 13.** *There exists a global certification scheme for BIPARTITENESS for n -vertex graphs whose identifiers are in the range $[M]$, with a certificate of size $O(n + \log \log M)$.*

Note that, in the standard case where M is polynomial in n , Corollary 13 gives a certificate of size $\Theta(n)$. Note also that this bound remains $\Theta(n)$ even in the case where $M = 2^{2^{O(n)}}$.

Proof of Theorem 12. Let us describe a global certification scheme for the existence of a homomorphism to H using a certificate of size $O(n \log n' + \log \log M)$ where $n' = |V(H)|$. First, since H has n' vertices, we can number these vertices from 1 to n' and write the index of a vertex of H on $\log n'$ bits. Similarly, for every $k, \ell \in \mathbb{N}$ with $k \leq \ell$, by applying Theorem 7, we can number the functions in $H_{k,\ell}$ from 1 and $|H_{k,\ell}|$. Thus, a function of $H_{k,\ell}$ can be represented using $\log |H_{k,\ell}| = O(k + \log \log \ell)$ bits.

Let $G = (V, E)$ be a graph with $|V| = n$, for which there exists a homomorphism φ from G to H . Let Id be an identifier assignment of G . The certificate given by the prover is the following one. Let us denote by $S := \{\text{Id}(v) \mid v \in V\}$ the set of identifiers appearing in G . The set S is included in $[M]$ and has size n . Let $h \in H_{n,M}$ be a perfect hash function for S .

By definition, the function h induces a bijection between S and $[n]$. Let L be the list of size n such that the i -th element of L , denoted by $L[i]$, is equal to $\varphi(v)$, where v is the unique vertex in V such that $h(\text{Id}(v)) = i$. The certificate given by the prover to the vertices is the triple (n, h, L) , where h is represented by its numbering in $H_{n,M}$. Since it uses $O(n \log n')$ bits to represent L and $O(n + \log \log M)$ bits to represent h , the overall size of the certificate is $O(n \log n' + \log \log M)$.

Let us describe the verification algorithm. Each vertex u does the following. First, it reads n in the global certificate and computes M . Then, it can determine h in $H_{n,M}$ thanks to its numbering in the certificate. Finally, u accepts if and only if, for all $v \in N(u)$, $\{L[h(\text{Id}(u))], L[h(\text{Id}(v))]\} \in E'$. If it is not the case, u rejects.

Let us prove the correctness. First, assume that G indeed admits a homomorphism to H . Then, by giving the certificate as described above, since φ is a homomorphism, each vertex $u \in V$ accepts. Conversely, assume that every vertex accepts with some certificate c , and let us prove that there exists a homomorphism from G to H . Since all the vertices accept, every vertex u checked if $\{L[h(\text{Id}(u))], L[h(\text{Id}(v))]\} \in E'$ for every $v \in N(u)$, for some function h which is written in c . Note that nothing ensures that h is indeed a perfect hash function for the set S of identifiers, but in fact, it is not necessary to check that h is injective on S . Indeed, since every vertex u accepted, then for every $v \in N(u)$, we have $\{L[h(\text{Id}(u))], L[h(\text{Id}(v))]\} \in E'$. So $\varphi(u) := L[h(\text{Id}(u))]$ defines a homomorphism from G to H . Thus, it proves the correctness of the scheme. ◀

We note that the idea of using perfect hashing has independently been used in [4] with another type of labeling, but to our knowledge, it is the first time that perfect hashing is used in distributed computing. We hope that this technique could have other applications in future works, in particular for problems related to space complexity.

References

- 1 Nicolas Bousquet, Laurent Feuilloley, and Sébastien Zeitoun. Brief announcement: Global certification via perfect hashing. In Ran Gelles, Dennis Olivetti, and Petr Kuznetsov, editors, *Proceedings of the 43rd ACM Symposium on Principles of Distributed Computing, PODC 2024*, pages 429–431. ACM, 2024. doi:10.1145/3662158.3662781.
- 2 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.
- 3 Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In Magnús M. Halldórsson and Shlomi Dolev, editors, *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15-18, 2014*, pages 367–376. ACM, 2014. doi:10.1145/2611462.2611493.
- 4 Louis Esperet, Nathaniel Harms, and Viktor Zamaraev. Optimal adjacency labels for subgraphs of cartesian products. *SIAM J. Discret. Math.*, 38(3):2181–2193, 2024. doi:10.1137/23M1587713.
- 5 Laurent Feuilloley. Introduction to local certification. *Discret. Math. Theor. Comput. Sci.*, 23(3), 2021. doi:10.46298/DMTCS.6280.
- 6 Laurent Feuilloley and Juho Hirvonen. Local verification of global proofs. In *32nd International Symposium on Distributed Computing, DISC 2018, New Orleans, LA, USA, October 15-19, 2018*, volume 121, pages 25:1–25:17, 2018. doi:10.4230/LIPIcs.DISC.2018.25.
- 7 Amos Korman, Shay Kutten, and David Peleg. Proof labeling schemes. *Distributed Comput.*, 22(4):215–233, 2010. doi:10.1007/S00446-010-0095-3.
- 8 Kurt Mehlhorn. *Data Structures and Algorithms 1: Sorting and Searching*, volume 1 of *EATCS Monographs on Theoretical Computer Science*. Springer, 1984. doi:10.1007/978-3-642-69672-5.

12 Renaming in distributed certification

- 9 Moni Naor, Merav Parter, and Eylon Yogev. The power of distributed verifiers in interactive proofs. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 1096–1115. SIAM, 2020. doi:10.1137/1.9781611975994.67.