Holographic Trees

Paolo Boldi[1] and Sebastiano Vigna[1*]

Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano, Via Comelico 39/41, I-20135 Milano MI, Italy. {vigna,boldi}@dsi.unimi.it

Abstract It is known that computations of anonymous networks can be reduced to the construction of a certain graph, the *minimum base* of the network. The crucial step of this construction is the inference of the minimum base from a finite tree that each processor can build (its *truncated view*). We isolate those trees that make this inference possible, and call them *holographic*. Intuitively, a tree is holographic if it is enough self-similar to be uniquely extendible to an infinite tree. This possibility depends on a *size function* for the class of graphs under examination, which we call a *holographic bound* for the class. Holographic bounds give immediately, for instance, bounds for the quiescence time of self-stabilizing protocols. In this paper we give weakly tight holographic bounds for some classes of graphs.

1 Introduction

This paper investigates combinatorial properties of trees and graphs whose very definition has been inspired by some problems in the study of distributed anonymous and self-stabilizing (synchronous) computations. In particular, we shall define and study *holographic trees* and *holographic bounds*, which turn out to play a major rôle in the construction of distributed anonymous algorithms and self-stabilizing protocols.

The reader might be easily bewildered by the amount of notions that must be absorbed to grasp the concepts above, and by the appearent opaqueness of Definition 2 and 3. Indeed, some knowledge of anonymous computations and theory of graph fibrations is necessary to understand them completely. Thus, the rest of this introduction is devoted to presenting a "historical" reconstruction of the notion of holographic tree and holographic bound, so to place these concepts in a computational perspective.

Consider a *network* of processors, represented by a strongly connected graph [1]. We assume that the network is *anonymous* [1, 13, 2], that is, all processors start from the same initial state and apply the same algorithm. The network is *synchronous*, in the sense that all processors take a step at the same time, and the new state of a processor depends on its own state and on the states of its in-neighbours.

One of the main concerns in the theory of anonymous computation is to establish which problems can be solved on a network [2]. This apparently gigantic task was enormously simplified by the discovery that the state of a processor after the k-th step of *any*

* The authors have been partially supported by the Italian MURST (Finanziamento di iniziative di ricerca "diffusa" condotte da parte di giovani ricercatori).

[1] Our graphs are directed, and may possess multiple arcs and loops—see Sect. 2.

[2] Or on a class of networks, that is, using an algorithm that will work for *every* network out of the class. The class is used to represent the *knowledge* processors possess about the network.

anonymous computation depends only on a finite tree, which is the truncation at depth k of an infinite tree, the *view* of the processor. In the bidirectional case, Yamashita and Kameda [13], inspired by the seminal work of Angluin[1], showed that views correspond to a standard graph-theoretic construction, the *universal cover*. Subsequently, it was shown [3] that in the general case views are the *universal total graphs* (in the sense of the theory of graph fibrations [7]) of the processors.

Let G be a network and i a processor of G. The view of G at i, denoted by $\widetilde{G}^i$, has the (finite) paths of G ending into i as nodes (the root of $\widetilde{G}^i$ being the empty path), with an arc from node π to node π' if π is obtained by adding an arc at the beginning of π'. (The reader may want to consult Fig. 2 to see an example of a view.)

It is not difficult to become convinced that each processor can anonymously build its own view truncated at any desired depth: at step k each processor gathers the view truncated at depth k from each of its in-neighbours, and it is thus able to build its own view truncated at depth $k+1$. An example of the first steps of this algorithm is given (for the network of Fig. 2) in Table 1 (for the time being, just ignore the last column).

The goal of each processor would be (in a perfect world) the computation of G and of the node of G that corresponds to the processor itself. However, this is in general impossible, since many processors can possess the same view, and thus cannot reach different states, no matter which algorithm they use. A more feasible goal, which can be indeed achieved, is the computation of the *minimum base* of the network, which is essentially[3] obtained by collapsing those processors that have the same view. (Again, an example can be found in Fig. 2.) Indeed, *every* anonymous computation can be factored into the computation of the minimum base followed by a local computation [4] [6].

The fundamental question we have to answer to is now: How deep must a truncation of the view be for the correct computation of the minimum base? In other words: How long must the algorithm sketched above run to compute a correct result?

The first answers were given originally by Yamashita and Kameda: if a bound N on the number of processors of the network is known, N^2 steps suffice. Subsequently, using a result of Norris [12], the bound was improved to $2N-1$, and eventually to $N+D$ (where D is an upper bound on the diameter) in [3].

Once the minimum base B has been computed on the basis of a certain truncated view T, one notes that T is a prefix[5] of a unique view of B. In fact, this infinite tree is the view of one or more processors of the network (among which the processor that computed T).

In other words, we can produce an infinite extension of T, in a way that respects the internal similarities of T, and this infinite extension is entirely described by a graph (the minimum base) with a selected node. With a bit of imagination, we shall call a tree for which this operation is possible (preserving some uniqueness properties) *holographic*, since, as in a hologram, it is a small piece that contains all the information of a much larger (indeed, infinite) picture. For instance, it is reasonable to think that sufficiently

[3] A more precise definition, which requires the introduction of graph fibrations, is given in Sect. 3.

[4] As long as a bound on the network size is known; if no bound on the size is known, this is not possible, and a completely different approach is necessary (see, e.g., [5]).

[5] For a formal definition of tree prefix, see Sect. 2.

deep truncated views will always be holographic. And this is indeed the reason why anonymous computations work: the bounds quoted above quantify (rather roughly, as we shall see) the depth after which a tree becomes holographic—in our terminology, they are *holographic bounds*.

Of course, holographic bounds depend on the class of networks under examination. If the class is small, very few levels (even just one, as shown in an example) will suffice to identify the minimum base. Large classes, instead, will require more information, and thus more levels (see, e.g., the lower bounds given in Sect. 4).

We now must admit that we partially lied to our reader: even if holographic bounds do provide upper bounds on the runs of anonymous computation, their real application lies in the domain of *self-stabilization*. A system is self-stabilizing if, for every initial state, after a finite number of steps it cannot deviate from a specified behaviour. Self-stabilization of distributed systems was introduced by Dijkstra in his celebrated paper [11], and has since become an important framework for the study of fault-tolerant computations.

In [4, 8] we showed the existence of a self-stabilizing protocol that computes the truncated view; essentially, after $D+1$ steps all processors possess a truncation of their view at least $D+1$ levels deep, no matter which the initial state was. Similarly to the anonymous case, this allows self-stabilization to arbitrary behaviours (for which a protocol exists) which depend only on the minimum base; thus, the quiescence time (i.e., the number of steps after which the desired behaviour starts) will depend on the number of levels required to compute the minimum base correctly—again, on a holographic bound for the class of networks under examination.

In this paper we show that the function assigning to a graph G the size $n_{\hat{G}} + D_{\hat{G}}$, where $\hat{G}$ is the minimum base of G, is a holographic bound for the class of all strongly connected graphs, and that this bound is weakly tight (there are infinite graphs for which the size must be at least $n_{\hat{G}} + D_{\hat{G}}$). We also show that the bound drops to $D_G + 1$ if the nodes are labelled injectively (i.e., in the network interpretation, if processors have unique identifiers). Also in this case, we provide matching lower bounds.

2 Graph-theoretical Definitions

A *(directed) (multi)graph* G is given by a nonempty set $N_G = \{1, 2, \dots, n_G\}$ of nodes and a set A_G of arcs, and by two functions $s_G, t_G : A_G \to N_G$ that specify the source and the target of each arc. A *(arc- and node-)coloured graph* (with set of colours C) is a graph endowed with a colouring function $\gamma : N_G + A_G \to C$ (the symbol $+$ denotes the disjoint union). We write $i \stackrel{a}{\to} j$ when the arc a has source i and target j, and $i \to j$ when $i \stackrel{a}{\to} j$ for some $a \in A_G$. We denote with D_G the diameter of G. Subscripts will be dropped whenever no confusion is possible.

A (in-directed) *tree* is a graph[6] with a selected node, the root, such that every node has exactly one directed path to the root. If T is a tree, we write $h(T)$ for its *height* (the length of the longest directed path). In every tree we consider in this paper, all maximal

[6] Since we need to manage infinite trees too, we allow the node set of a tree to be $\mathbf{N}$.

paths have length equal to the height. We write $T \upharpoonright k$ for the tree T truncated at height k, that is, we eliminate all nodes at distance greater than k from the root.

Trees are partially ordered[7] by *prefix*, that is, $T \leq U$ iff $T \cong U \upharpoonright h(T)$; this partial order is augmented with a bottom element $\perp$, with $h(\perp) = -1$ by definition (so h is strictly monotonic). The infimum in this partial order, denoted by $\wedge$, is the tallest common prefix (or $\perp$ if no common prefix exists). The supremum between T and U exists iff T and U are comparable.

3 Graph Fibrations

In this paper we exploit the notion of *graph fibration* [7]. A fibration, which is essentially a local in-isomorphism, formalizes the idea that processors that are connected by arcs with the same colours to processors behaving in the same way (with respect to the colours) will behave alike. In this section we gather (without proof) a number of definitions and results about graph fibrations; although some of the statements are true for all graphs, for sake of simplicity we shall assume that all graphs (except for trees) are strongly connected.

Recall that a *graph morphism* $f : G \to H$ is given by a pair of functions $f_N : N_G \to N_H$ and $f_A : A_G \to A_H$ that commute with the source and target functions, that is, $s_H \circ f_A = f_N \circ s_G$ and $t_H \circ f_A = f_N \circ t_G$. (The subscripts will usually be dropped.) In other words, a morphism maps nodes to nodes and arcs to arcs in such a way to preserve the incidence relation. Colours on nodes and arcs must be preserved too.

Definition 1. *A* fibration[8] *between (coloured) graphs G and B is a morphism $\varphi : G \to B$ such that for each arc $a \in A_B$ and for each node $i \in N_G$ satisfying $\varphi(i) = t(a)$ there is a unique arc $\widetilde{a}^i \in A_G$ (called the* lifting of a at i*) such that $\varphi(\widetilde{a}^i) = a$ and $t(\widetilde{a}^i) = i$.*

We recall some topological terminology. If $\varphi : G \to B$ is a fibration, G is called the *total graph* and B the *base* of φ. We shall also say that G is *fibred (over B)*. The *fibre* over a node $i \in N_B$ is the set of nodes of G that are mapped to i, and will be denoted by $\varphi^{-1}(i)$.

In Fig. 1 we sketched a fibration between two graphs. Note that, because of the lifting property described in Definition 1, all black nodes have exactly two incoming arcs, one (the dotted arc) going out of a white node, and one (the continuous arc) going out of a grey node. In other words, the in-neighbour structure of all black nodes *is the same*.

There is a very intuitive characterization of fibrations based on the concept of local isomorphism. A fibration $\varphi : G \to B$ induces an equivalence relation between the nodes of G, whose classes are precisely the fibres of φ. When two nodes i and j are equivalent (i.e., they are in the same fibre), there is a bijective correspondence between

7 We are in fact considering trees up to isomorphism (technically $\leq$ is just a preorder).

8 The name "fibration" comes from the categorical and homotopical tradition; indeed, our elementary definition is simply a restatement of the condition that $\varphi : G \to B$ induces a functor that is a fibration [9] between the free categories generated by G and B.

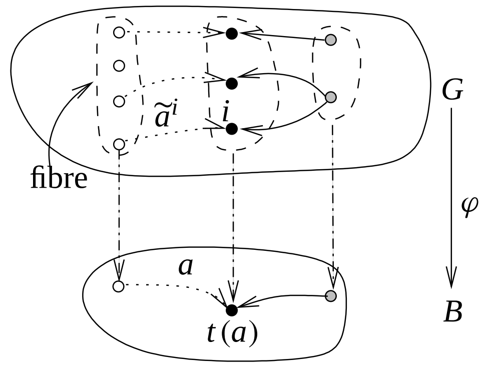

Figure1. A fibration.

arcs coming into i and arcs coming into j that preserves colours, and such that the sources of any two corresponding arcs are equivalent.

Let now G be a graph and i a node of G. We define an in-directed rooted coloured tree $\widetilde{G}^i$ as follows:

- the nodes of $\widetilde{G}^i$ are the (finite) paths of G ending in i, the root of $\widetilde{G}^i$ being the empty path; each node is given the same colour as the starting node of the path;
- there is an arc from the node π to the node π' if π is obtained by adding an arc a at the beginning of π' (the arc will have the same colour as a).

The tree $\widetilde{G}^i$ (which is always infinite if G is strongly connected and with at least one arc) is called the *universal total graph of G at i*, or, following Yamashita and Kameda, the *view of i*. We can define a graph morphism υ_G^i from $\widetilde{G}^i$ to G, by mapping each node π of $\widetilde{G}^i$ (i.e., each path of G ending in i) to its starting node, and each arc of $\widetilde{G}^i$ to the corresponding arc of G. The following important property holds:

Lemma 1. *For every node i of a graph G, the morphism $\upsilon_G^i : \widetilde{G}^i \to G$ is a fibration, called the* universal fibration of G at i.

The view at i is a tree representing intuitively "everything processor i can learn from interaction with its neighbours in an anonymous computation"; it plays in the general case the same rôle as the universal covering in the undirected case [13]. As we remarked in the introduction, it is not difficult to see that each processor can anonymously build its own view truncated at any desired depth (and k steps are needed to obtain k levels). Note that views of finite (strongly connected) graphs are exactly the *regular trees without leaves*, in the sense of Courcelle [10].

Consider now a graph B such that every fibration with total graph B is an isomorphism. Such a graph is called *fibration prime*: intuitively, fibration prime graphs cannot be "collapsed" by a fibration. It is worth observing that they are node rigid (i.e., all automorphisms act on the nodes as the identity), so, in particular, if B and B' are isomorphic and fibration prime then all isomorphisms $B \to B'$ coincide on the nodes. We have the following

Lemma 2. *Let $\varphi : G \to B$ and $\varphi' : G \to B'$ be fibrations, with B and B' fibration prime. Then $B \cong B'$.*

In other words, to each graph G we can associate a fibration prime graph $\hat{G}$, the *minimum base* of G, and a *minimal fibration* $\mu_G : G \to \hat{G}$ (in fact, there are several candidates for μ_G, but they are all defined in the same way on the nodes, and in this paper we shall use only the node component of minimal fibrations). In Fig. 2, we illustrate these notions by showing a graph G, its minimum base $\hat{G}$ and one of its views, $\widetilde{G}^1$. (The numbers shown on G are actual node names, and not colours; the numberings on $\hat{G}$ and $\widetilde{G}^1$ illustrate μ_G and υ_G^1.) There are three important comments to be made

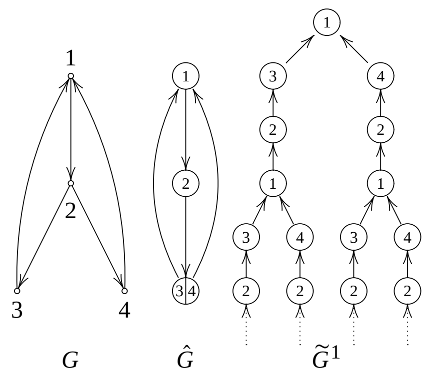

Figure2. A graph, its minimum base and a view.

about $\hat{G}$:

- fibration prime graphs (in particular, $\hat{G}$) have distinct views (i.e., $\widetilde{B}^i \cong \widetilde{B}^j$ iff $i = j$);
- $\hat{G}$ can be constructed by identifying isomorphic subtrees of $\widetilde{G}^i$ (the choice of $i \in N_G$ is irrelevant), and μ_G maps node i to the equivalence class containing $\widetilde{G}^i$;
- $\widetilde{G}^i \cong \widetilde{\hat{G}}^{\mu_G(i)}$, so we can compute $\mu_G(i)$ by searching for the node of $\hat{G}$ having $\widetilde{G}^i$ as view.

The fundamental fact we shall use intensively in all proofs is that the above considerations, which involve infinite objects, can be described by means of finite entities using the theorems of Sect. 5.

4 Holographic Bounds

Armed with our basic definitions, we now introduce the main concept we shall deal with. Our interest is in isolating those trees that are enough coherent, and contain enough information, to "replicate" themselves *ad infinitum.*

Definition 2. *Let $\mathscr{C}$ be a class of graphs. A* size function *for $\mathscr{C}$ is a map $\nu : \mathscr{C} \to \mathbf{N}$. Given a class $\mathscr{C}$, a size function ν and a finite tree T we define*

$$\mathscr{U}_{\mathscr{C},\nu}(T) = \left\{ \langle \hat{H}, \mu_H(j) \rangle \;\middle|\; H \in \mathscr{C},\, j \in N_H,\, \nu(H) \le h(T) \text{ and } T \le \widetilde{H}^j \right\}.$$

*We say that T is ($\langle \mathscr{C}, \nu \rangle$-)*holographic *iff $\mathscr{U}_{\mathscr{C},\nu}(T)$ is nonempty, and for every $\langle B, i \rangle, \langle B', i' \rangle \in \mathscr{U}_{\mathscr{C},\nu}(T)$ there is an isomorphism $\alpha : B \to B'$ such that $\alpha(i) = i'$ (we shall often state this condition by saying that $\mathscr{U}_{\mathscr{C},\nu}(T)$ contains* essentially *one element).*

The idea behind the definition above is that the set $\mathscr{U}_{\mathscr{C},\nu}(T)$ contains all the possible candidates for the pointed minimum bases of the graph (network) that generated T. The height of T is also used to confine the search to those graphs whose size is not too large. A holographic tree is a tree that is sufficiently self-similar to identify a single candidate (up to isomorphism).

Definition 3. *We say that ν is a* holographic bound *for $\mathscr{C}$ if for all $G \in \mathscr{C}$, all $i \in N_G$ and all $k \ge \nu(G)$ we have that $\widetilde{G}^i \restriction k$ is $\langle \mathscr{C}, \nu \rangle$-holographic.*

Note that the set $\mathscr{U}_{\mathscr{C},\nu}\left(\widetilde{G}^i \restriction k\right)$ is *always* nonempty, as it contains $\langle \hat{G}, \mu_G(i) \rangle$. A small holographic bound makes holographic more trees, as it lowers the height required by Definition 3, and moreover reduces the number of candidates; thus, given a class of graphs one is interested in finding out a small holographic bound. Indeed, such a bound provides also an upper bound on the quiescence time of self-stabilizing protocols [4, 8] (or on the running time of anonymous algorithms) for the class under consideration.

As an example, in Table 1 we display the first steps of the execution of the standard view construction algorithm on the network of Fig. 2, where T_i is $\widetilde{G}^i \restriction t$, that is, the tree constructed by processor i at time t. The last column gives, at each step, the content of $\mathscr{U}_{\mathscr{G},\nu}(T_2)$ (with $\nu(G) = n_{\hat{G}} + D_{\hat{G}}$, and $\mathscr{G}$ the class of all strongly connected graphs). The reader may notice that the second processor "changes his mind" a few times about $\hat{G}$, but ultimately its guess is correct.

We now prove a very general and intuitive property of holographic bounds: a holographic bound for $\mathscr{C}$ works for every subclass of $\mathscr{C}$, and moreover size functions pointwise larger than a holographic bound are still holographic bounds. Formally,

Theorem 1. *Let ν be a holographic bound for a class $\mathscr{C}$. Then every size function $\nu' \ge \nu$ is a holographic bound for every class $\mathscr{C}' \subseteq \mathscr{C}$.*

Proof. Let G be a graph of $\mathscr{C}$, i a node of G and $k \ge \nu'(G)$. We have to prove that $\mathscr{U}_{\mathscr{C}',\nu'}\left(\widetilde{G}^i \restriction k\right)$ contains essentially one element. But since, as it is immediate to check from Definition 2,

$$\mathscr{U}_{\mathscr{C}',\nu'}(T) \subseteq \mathscr{U}_{\mathscr{C},\nu}(T),$$

and the right-hand set contains essentially one element when $T = \widetilde{G}^i \restriction k$, we have the thesis. ◻

5 A Holographic Bound for All Graphs

In this section we provide a holographic bound for the class of all (strongly connected) graphs (note that by assuming $|C| = 1$, the bound applies also to noncoloured graphs).

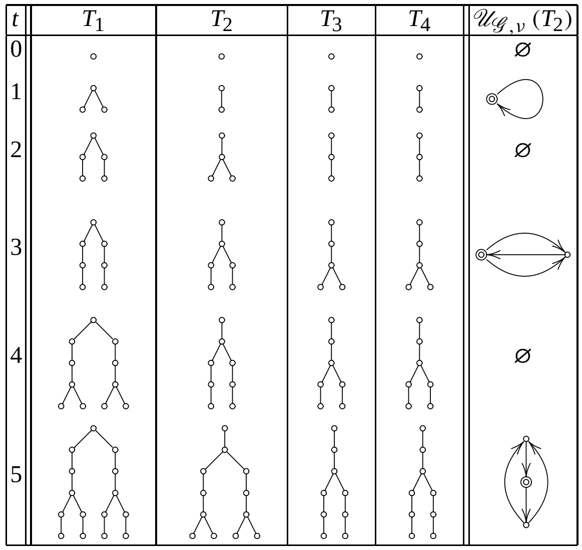

Table1. The first steps of the view construction algorithm on the network G of Fig. 2.

The bound is based on a number of graph-theoretical results, which show that sufficiently deep truncations of views characterize the (minimum bases of the) graph that generated them. First of all, we recall a result of Norris [12]:

Theorem 2. *$\widetilde{G}^i \cong \widetilde{G}^j$ iff $h\left(\widetilde{G}^i \wedge \widetilde{G}^j\right) \geq n-1$.*

Using the previous theorem, we prove the following

Theorem 3. *Let G be a strongly connected graph and B a fibration prime graph with minimum number of nodes satisfying $h\left(\widetilde{G}^i \wedge \widetilde{B}^j\right) \geq n_G + D_G$ for some $i \in N_G$ and $j \in N_B$: then there is a (minimal) fibration $\varphi : G \to B$ such that $\varphi(i) = j$; in particular, $B \cong \hat{G}$.*

Proof. Note that B has at most n nodes, because the minimum base of G satisfies the hypotheses. We shall build a morphism $\varphi : G \to B$ by sending a node l of G to the unique node $\varphi(l)$ of B satisfying $\widetilde{G}^l \upharpoonright (n-1) \cong \widetilde{B}^{\varphi(l)} \upharpoonright (n-1)$. This node can be found as follows: there is certainly a node $l' \in (\upsilon_G^i)^{-1}(l)$ which is at depth D at most. Thus, the subtree under l' in $\widetilde{G}^i \upharpoonright n+D$ has height at least $n-1$. Let $\psi : \widetilde{G}^i \upharpoonright (n+D) \to \widetilde{B}^j \upharpoonright (n+D)$ be the isomorphism above. Then $\varphi(l) = (\upsilon_B^j \circ \psi)(l')$. Note that the choice of l' is irrelevant, by the considerations about the views of fibration prime graphs made in Sect. 3.

We now define analogously φ on the arcs, by using the lifting property. Let a be an arc of G. We choose, as before, $l \in (\upsilon_G^i)^{-1}(t(a))$ which is at depth D at most, and consider the lifting $\widetilde{a}^l$. Then we set $\varphi(a) = (\upsilon_B^j \circ \psi)(\widetilde{a}^l)$. Note that this is compatible with our definition on the nodes, because $s(\widetilde{a}^l)$ is at depth $D+1$ at most, and thus its image through $\upsilon_B^j \circ \psi$ must be $\varphi(s(a))$, by Theorem 2. It is then easy to check that since

φ has been defined by a lifting and composition with isomorphisms and fibrations, it is itself a fibration. Moreover, by its very definition it maps i to j. $\square$

Corollary 1. *Let B_1 and B_2 be fibration prime, i_1 a node of B_1 and i_2 a node of B_2. If*

$$h\left(\widetilde{B_1}^{i_1} \wedge \widetilde{B_2}^{i_2}\right) \geq \max\{n_{B_1} + D_{B_1}, n_{B_2} + D_{B_2}\}$$

then there is an isomorphism $\alpha : B_1 \to B_2$ such that $\alpha(i_1) = i_2$.

The previous result shows that fibration prime graphs sharing enough levels of one of their views are isomorphic (and the nodes to which the view are associated are in correspondence). This is all we need to prove our first holographic bound:

Theorem 4. *The function mapping a graph G to $n_{\hat{G}} + D_{\hat{G}}$ is a holographic bound for the class of all graphs (hence for every class).*

Proof. By definition, we have to show that for every graph G, every node i of G and every $k \geq n_{\hat{G}} + D_{\hat{G}}$ the class

$$\left\{ \langle \hat{H}, \mu_H(j) \rangle \;\middle|\; j \in N_H, n_{\hat{H}} + D_{\hat{H}} \leq k \text{ and } \widetilde{G}^i \upharpoonright k \leq \widetilde{H}^j \right\}$$

contains essentially one element. To this purpose, we take an arbitrary pair $\langle B, l \rangle$ from the set and show how to build an isomorphism between $\hat{G}$ and B that maps $\mu_G(i)$ to l. But since $\widetilde{G}^i \upharpoonright k \leq \widetilde{B}^l$ and $k \geq \max\left\{n_B + D_B, n_{\hat{G}} + D_{\hat{G}}\right\}$, we can apply Corollary 1 to $\hat{G}$ and B, with chosen nodes $\mu_G(i)$ and l (recall that $\widetilde{G}^i \cong \widetilde{\hat{G}}^{\mu_G(i)}$). $\square$

An example of trees that are holographic using the bound above is given in Table 2, where we show all the (uncoloured) trees of height at most four and indegree at most two that are holographic, together with the base and the node that generated them. One should compare this list with some of the nonholographic trees appearing in the execution of the view construction algorithm in Table 1. The reader might now be curious to know whether $n_{\hat{G}} + D_{\hat{G}}$ is the best possible holographic bound. We do not know the answer to this question, but we have some partial results. Let $\mathscr{C}$ be any class of graph including the fibration prime graphs shown in Fig. 3. Note that both $G_{n,D}$ and $H_{n,D}$ have n nodes and diameter D (the difference between the two families is given by the positioning of the dotted arc). It is tedious but easy to check that for all D and n we have that $\widetilde{G_{n,D}}^1$ and $\widetilde{H_{n,D}}^1$ are isomorphic up to level $n + D - 1$, but not up to level $n + D$, and this property is the key to the proof of the following lower bounds:

Theorem 5. *Every holographic bound for a class $\mathscr{C}$ as above is at least $n_{\hat{G}} + D_{\hat{G}}$ for an infinite number of graphs.*

Proof. By contradiction, suppose that $\nu(G) \geq n_{\hat{G}} + D_{\hat{G}}$ holds only for finitely many $G \in \mathscr{C}$. Thus, there are n and D such that $\nu(G_{n,D}), \nu(H_{n,D}) < n + D$. But then $\mathscr{U}_{\mathscr{C},\nu}\left(\widetilde{G_{n,D}}^1 \upharpoonright n + D - 1\right)$ contains both $\langle G_{n,D}, 1 \rangle$ and $\langle H_{n,D}, 1 \rangle$, a contradiction. $\square$

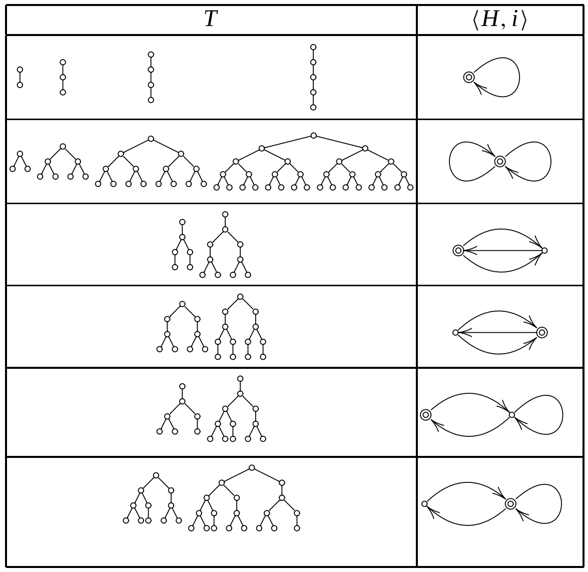

Table2. Small holographic trees.

Theorem 6. *A class $\mathscr{C}$ as above has no holographic bound depending only on the number of nodes and on the diameter (or on the number of nodes and on the diameter of the minimum base) that is smaller than $n_{\hat{G}} + D_{\hat{G}}$ for some $G \in \mathscr{C}$.*

Proof. Let $L \in \mathscr{C}$ be a graph with n nodes and diameter D such that $\nu(L) < n_{\hat{L}} + D_{\hat{L}}$; in the first case, $\nu(G_{n,D}) = \nu(H_{n,D}) = \nu(L) < n_{\hat{L}} + D_{\hat{L}} \leq n + D$, and we proceed as in the proof of Theorem 5. In the second case, we have $\nu\left(G_{n_{\hat{L}},D_{\hat{L}}}\right) = \nu\left(H_{n_{\hat{L}},D_{\hat{L}}}\right) = \nu(L) < n_{\hat{L}} + D_{\hat{L}}$, and the thesis follows again. □

In the rest of the section, we highlight two worked out examples specializing Theorem 4.

Inregular graphs. Since the inregular graphs are exactly the total graphs over bouquets (i.e., graphs with exactly one node), we have that the constant function 1 is the (obviously minimum) holographic bound for that class, and the holographic trees are exactly the inregular trees with at least one arc.

Complete multipartite graphs. A graph is *complete multipartite* iff its node set can be partitioned into independent[9] sets, and there is exactly one arc from node i to node j when i and j do not belong to the same part. The minimum base of a complete multipartite graph G can be constructed as follows: let $k_1, k_2, \ldots, k_{l_G}$ be a list (without repetitions) of the cardinalities of the parts of G, and $m_1, m_2, \ldots, m_{l_G}$ be the respective multiplicities (i.e., m_i is the number of parts of cardinality k_i). The graph $\hat{G}$ has l_G

[9] A set of nodes is *independent* iff there are no arcs with source and target in the set.

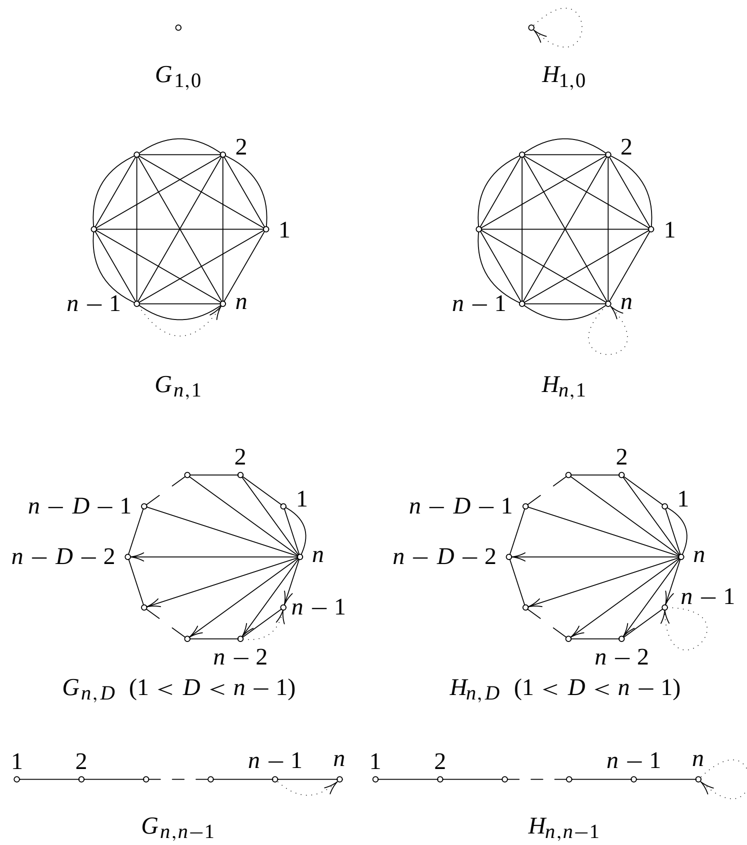

Figure3. Graphs with similar views.

nodes, and the number of arcs from node i to node j is $m_i k_i$ if $i \neq j$, $(m_i - 1)k_i$ otherwise. As a consequence, a holographic bound for this class is given by $\nu(G) = l_G + 1$.

6 A Holographic Bound for All Labelled Graphs

In this section we prove that $D_G + 1$ is a holographic bound for the class $\mathscr{L}$ of labelled graphs, that is, graphs whose nodes are coloured injectively. All such graphs are obviously fibration prime. Note that the theory described in Sect. 3 specializes, and that also views carry on the colouring (which however is no longer injective). First, we show that Theorem 3 can be restated as follows:

Theorem 7. *Let G and B be labelled graphs, and suppose $h\left(\widetilde{G}^i \wedge \widetilde{B}^j\right) \geq D_G + 1$ for some $i \in N_G$ and $j \in N_B$: then there is an isomorphism $\alpha : G \cong B$, and $\alpha(i) = j$.*

Proof. We define α by sending a node of G to the unique node of B having the same colour, and on the arcs by lifting. More in detail, let a be an arc of G. We choose a $l \in (\upsilon_G^i)^{-1}(t(a))$ which is at depth D at most, and consider the lifting $\widetilde{a}^l$. Then we set $\alpha(a) = (\upsilon_B^j \circ \psi)(\widetilde{a}^l)$, where ψ is the isomorphism from $\widetilde{G}^i \restriction D_G + 1$ to $\widetilde{B}^j \restriction D_G + 1$. Note that this is compatible with our definition on the nodes, because $s(\widetilde{a}^l)$ is at depth $D + 1$ at most, and thus its image through $\upsilon_B^j \circ \psi$ must have the same colour as $s(a)$. It is then easy to check that since α has been defined by a lifting and composition with isomorphisms and fibrations, it is itself a fibration, and thus an isomorphism by primality of G and B. Moreover, by its very definition it maps i to j. ◻

Theorem 8. *The function mapping a graph G to $D_G + 1$ is a holographic bound for the class of all labelled graphs (hence for every class of labelled graphs).*

Proof. By definition, we have to show that for every labelled graph G, every node i of G and every $k \geq D_G + 1$ the class

$$\left\{ \langle H, j \rangle \;\middle|\; H \text{ is labelled, } j \in N_H,\ D_H + 1 \leq k \text{ and } \widetilde{G}^i \restriction k \leq \widetilde{H}^j \right\}$$

contains essentially one element (we omitted the hat symbols as all labelled graphs are fibration prime). But if $\langle H, j \rangle$ is an arbitrary element from the set, we have $h\left(\widetilde{G}^i \wedge \widetilde{H}^j\right) \geq k \geq D_G + 1$, and by Theorem 7 we obtain the thesis. ◻

Finally, as in the previous section, we prove some lower bounds:

Theorem 9. *Every holographic bound ν for $\mathscr{L}$ is at least $D + 1$ for an infinite number of graphs.*

Proof. Suppose $\nu(G) < D_G + 1$, and consider nodes i and j of G which maximize the distance from j to i. Let H be the graph obtained from G by adding an additional loop at j. Clearly $\mathscr{U}_{\mathscr{L},\nu}\left(\widetilde{G}^i \restriction D_G\right)$ would contain $\langle H, i \rangle$, unless $\nu(H) \geq D_G + 1 = D_H + 1$. This gives immediately the result. ◻

An absolutely analogous proof shows also that

Theorem 10. *The class $\mathscr{L}$ has no holographic bound depending only on the diameter that is smaller than $D_G + 1$ for some $G \in \mathscr{L}$.*

References

[1] Dana Angluin. Local and global properties in networks of processors. In *Proc. 12th Symposium on the Theory of Computing*, pages 82–93, 1980.

[2] Paolo Boldi, Bruno Codenotti, Peter Gemmell, Shella Shammah, Janos Simon, and Sebastiano Vigna. Symmetry breaking in anonymous networks: Characterizations. In *Proc. 4th Israeli Symposium on Theory of Computing and Systems*, pages 16–26. IEEE Press, 1996.

[3] Paolo Boldi and Sebastiano Vigna. Computing vector functions on anonymous networks. In Danny Krizanc and Peter Widmayer, editors, *SIROCCO '97. Proc. 4th International Colloquium on Structural Information and Communication Complexity*, volume 1 of *Proceedings in Informatics*, pages 201–214. Carleton Scientific, 1997. An extended abstract appeared also as a Brief Announcement in *Proc. PODC '97*, ACM Press.

[4] Paolo Boldi and Sebastiano Vigna. Self-stabilizing universal algorithms. In Sukumar Ghosh and Ted Herman, editors, *Self–Stabilizing Systems (Proc. of the 3rd Workshop on Self–Stabilizing Systems, Santa Barbara, California, 1997)*, volume 7 of *International Informatics Series*, pages 141–156. Carleton University Press, 1997.

[5] Paolo Boldi and Sebastiano Vigna. Computing anonymously with arbitrary knowledge. In *Proc. 18th ACM Symposium on Principles of Distributed Computing*, pages 181–188. ACM Press, 1999.

[6] Paolo Boldi and Sebastiano Vigna. An effective characterization of computability in anonymous networks. In Jennifer L. Welch, editor, *Distributed Computing. 15th International Conference, DISC 2001*, number 2180 in Lecture Notes in Computer Science, pages 33–47. Springer–Verlag, 2001.

[7] Paolo Boldi and Sebastiano Vigna. Fibrations of graphs. *Discrete Math.*, 243:21–66, 2002.

[8] Paolo Boldi and Sebastiano Vigna. Universal dynamic synchronous self-stabilization. *Distr. Comput.*, 15, 2002.

[9] Francis Borceux. *Handbook of Categorical Algebra 2*, volume 51 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, 1994.

[10] Bruno Courcelle. Fundamental properties of infinite trees. *Theoret. Comput. Sci.*, 25(2):95–169, 1983.

[11] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *CACM*, 17(11):643–644, 1974.

[12] Nancy Norris. Universal covers of graphs: Isomorphism to depth $n-1$ implies isomorphism to all depths. *Discrete Appl. Math.*, 56:61–74, 1995.

[13] Masafumi Yamashita and Tiko Kameda. Computing on anonymous networks: Part I—characterizing the solvable cases. *IEEE Trans. Parallel and Distributed Systems*, 7(1):69–89, 1996.