

Fibrations of Graphs

Paolo Boldi*

Sebastiano Vigna*

Abstract

A *fibration* of graphs is a morphism that is a local isomorphism of in-neighbourhoods, much in the same way a covering projection is a local isomorphism of neighbourhoods. This paper develops systematically the theory of graph fibrations, emphasizing in particular those results that recently found application in the theory of distributed systems.

Keywords: graph fibrations, graph coverings, graph factorizations.

Contents

1	Introduction	2
2	Definitions and basic properties	3
2.1	Graph-theoretical definitions	3
2.2	Fibrations and coverings	4
2.3	Groups, fibrations and automorphisms	7
2.4	Some properties of fibrations between finite graphs	9
2.5	An application	10
3	Universal fibrations and coverings	11
3.1	Universal total graphs	11
3.2	Universal coverings	12
3.3	Nodes with the same universal total graph	13
4	Minimal fibrations	15
4.1	Constructing minimum bases	18
5	Graphs fibred over bouquets	19
5.1	Factorization lemmata	21
5.2	Regular Graphs	24
5.3	Schreier Graphs	25
6	A categorical standpoint	27
6.1	Pullbacks	29
6.2	The category of fibrations over a given base	32
6.3	Counting minimal fibrations of the cycle	34
7	Open problems	36

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

1 Introduction

A morphism of (directed multi)graphs $\varphi : G \rightarrow B$ is a *fibration* when each arc of B can be uniquely lifted at every node in the fibre of its target. This simple definition implies that locally φ is an isomorphism of in-neighbourhoods. In this paper we develop the theory of graph fibrations, with a special emphasis on results related to some applications in computer science; we shall also pay attention to the mutual relations between fibrations, group actions and categorical constructions.

Historically, the definition of graph fibration can be traced back to the first papers about fibrations between categories, which were in turn inspired by the notion of fibration in homotopy theory. John Gray [13], in one of the oldest paper on the subject, attributes the definition to Alexandre Grothendieck [15, 1], who devised it at the end of the '50s in connection with his work on the foundations of algebraic geometry. A graph is just a “category without composition and identities”, and thus the definition of fibration between categories applies to graphs just by taking the free categories they generate (this point will be fully explained in Section 6). The definition we shall use is exactly an elementary restatement of the categorical notion.

To be true, the genealogy is a bit more involved. Independently, at the end of the '60s, Horst Sachs introduced the concept of *divisor* (*Teiler*) of a graph [29, 30, 25], which was intensively studied by the community working on algebraic graph theory (for a detailed description and bibliography, see Chapter 4 of [11]). In our terminology, a (strongly connected) graph B is a rear divisor of a graph G exactly when there is a fibration $\varphi : G \rightarrow B$, but this fibration is not part of the definition, and in general there are many different fibrations between a graph and one of its rear divisors (dually for front divisors and opfibrations). One could say that divisors are to fibrations as partitions of the integers are to functions between finite sets.

Divisors can be used to factor the characteristic polynomial of a graph (i.e., the characteristic polynomial of its adjacency matrix), as it can be shown that the characteristic polynomial of B divides the characteristic polynomial of G (and the quotient has integer coefficients). The categorical and graph-theoretical communities seem to have never been aware of the relation between the two concepts (of course, the connection between divisibility of *undirected* graphs and topological coverings was known at least since [29]).

In the early '70s, Allen Schwenk [32] (building on a result by Abbe Mowshowitz [22]) introduced the notion of *equitable partition* of the vertices of an undirected graph G , and showed that the characteristic polynomial of G is divisible by the characteristic polynomial of a certain matrix induced by the partition (for instance, the *degree partition* of G [19] is equitable). Indeed, such a matrix is the adjacency matrix of a front divisor of G , but Schwenk seems to be unaware of this fact (irony of fate, he just missed it, as he reviewed [31] Petersdorf and Sachs' next paper [26] on a connected subject).

A related research area, started by Tomaž Pisanski and Jože Vrabec [27], is concerned with the concept of *graph bundle*, a topological (as opposed to categorical/combinatorial) generalization of the notion of covering. Essentially, a graph bundle with base G and fibre F (both being undirected graphs) is the 1-skeleton of a topological locally trivial bundle over B with fibre F and structure group $\mathbf{Aut}(F)$. This kind of bundles are a particular case of (categorically defined) fibrations between *symmetric reflexive* graphs, as we shall briefly discuss in Section 6; in this case, the underlying combinatorial structure is very different from ours. It is interesting to note that by a mysterious coincidence Younki Chae, Jin Ho Kwak and Jaeun Lee [10], starting from Schwenk's work, studied the problem of computing the characteristic polynomial of a graph bundle, rediscovering in the case of a discrete fibre Sach's original result [30] that the characteristic polynomial of a graph divides the characteristic polynomial of its covering spaces.

Our main motivation for the study of graph fibrations comes from the theory of distributed

systems. In the early '80s, a seminal paper by Dana Angluin [2] introduced undirected graph coverings (and in particular, universal coverings) as a way for proving impossibility results on bidirectional *anonymous* networks (viz., networks where all processors are identical and start from the same state). The paper also posed a number of interesting mathematical questions, leading, for instance, to Frank Leighton's proof of Angluin's conjecture that every two undirected finite graphs with the same universal covering have a common finite cover [19]. Eventually, a complete characterization by means of undirected graph coverings was obtained by Masafumi Yamashita and Tiko Kameda [37, 36].

In the last years, it has been realized that graph coverings are no longer sufficient to solve analogous problems in a more general setting, that is, when the processors of a network are only able to transmit messages by broadcast or links are unidirectional. It has turned out that the right mathematical notion in this case is exactly that of fibration; indeed, fibrations have been used to solve completely problems such as leader election [5] and function computation [6] in general anonymous networks. In turn, these new applications have stimulated new research and created new mathematical problems, which are (at least partially) addressed in this paper.

We start by summarizing the graph-theoretical definitions and basic properties we are going to use. Then, in Section 3 we discuss universal fibrations and coverings. The theory of minimum bases is developed in Section 4, where we show that having the same minimum base is equivalent to having the same universal total graphs. Each section is completed by a short informal discussion of the related applications to distributed systems. Section 5 studies graphs fibred over bouquets. Finally, in Section 6 we develop a general categorical framework, showing in particular that fibrations are preserved by pullbacks. This allows to prove several theorems about common fibrations and coverings of graphs. Moreover, we give a representation theorem in terms of functor categories that allows one to study counting problems, and by way of example we count the number of nonisomorphic minimal fibrations of a bidirectional cycle.

A note is needed about the meaning of the word “graph” in this paper. We adopt Berge's point of view [4]: all graphs are directed, possibly infinite, and can possess loops and multiple arcs. When necessary, inside this larger class we single out *separated* graphs (that do not possess multiple arcs), *loopless* graphs, and so on. We also discuss undirected graphs under the form of *symmetric* graphs, that is, graphs with a *specified* involution on the arc set that exchanges source and target of each arc (unfortunately, the word “symmetric” has sometimes been used with different meanings in the graph-theoretical literature). Every undirected graph has a symmetric representation, but the converse is not true, as there are two kind of loops: the ones that are fixed by the symmetry and the ones that are not. There is no way of translating this difference in the language of undirected graphs, and this is probably the reason why loops have always been so disturbing in the study of coverings (a full discussion of this issue can be found at the end of Section 5).

We consider symmetry (not a property of but rather) a *structure* on a graph. As a consequence, morphisms between symmetric graphs must preserve symmetry. The definition of symmetric fibration and covering we use (which turn out to be equivalent) are naturally induced by this point of view. Note that, whenever fit, we shall draw an undirected edge in place of a pair of opposite directed arcs (even when the graph is not symmetric).

2 Definitions and basic properties

2.1 Graph-theoretical definitions

A (*directed multi*)graph G is defined by a set N_G of nodes and a set A_G of arcs, and by two functions $s_G, t_G : A_G \rightarrow N_G$ that specify the source and the target of each arc (we shall drop the

subscripts whenever no confusion is possible). We use the notation $G(x, y)$ for denoting the set of arcs from x to y , that is, the set of arcs $a \in A_G$ such that $s(a) = x$ and $t(a) = y$; the arcs in $G(x, y)$ are said to be *parallel* to one another. A *loop* is an arc with the same source and target. Following common usage, we denote with $G(-, x)$ the set of arcs coming into x , that is, the set of arcs $a \in A_G$ such that $t(a) = x$, and analogously with $G(x, -)$ the set of arcs going out of x . A graph is *locally finite* if $G(x, -)$ and $G(-, x)$ are finite for every node x .

A *symmetric graph* is a graph endowed with a symmetry, that is, an involution (a self-inverse bijection) $(\bar{\quad}) : A_G \rightarrow A_G$ such that $s(a) = t(\bar{a})$ (and consequently $t(a) = s(\bar{a})$) for all arcs $a \in A_G$. A *semi-edge* of a symmetric graph is a loop a such that $\bar{a} = a$. Given a graph G , we define its (formal) *symmetrization* $\mathbf{Sym}(G)$ as the graph obtained by adding for each arc $a \in G(x, y)$ a new arc \bar{a} going from y to x , with symmetry defined in the obvious way.

A graph G is *j -inregular* (*k -outregular*) if $|G(-, x)| = j$ ($|G(x, -)| = k$, respectively). A *j -inregular, k -outregular graph* is said to be *(j, k) -regular*. For finite or symmetric graphs *(j, k) -regularity implies $j = k$* , and when $j = k$ we simply say that G is *j -regular*.

A *path* (of length n) is a sequence $x_0 a_1 x_1 \cdots x_{n-1} a_n x_n$, where $x_i \in N_G$, $a_j \in A_G$, $s(a_j) = x_{j-1}$ and $t(a_j) = x_j$. We shall usually omit the nodes from the sequence when at least one arc is present. If G is symmetric, a path is called *symmetrically stuttering* (or, simply, *stuttering*) iff it contains a subpath of the form $a\bar{a}$; a *nonstuttering walk* of a graph G is a nonstuttering path of $\mathbf{Sym}(G)$. Since we shall only be concerned with walks of this kind, we shall drop the adjective “nonstuttering” in the sequel. We shall say that G is (strongly) *connected* iff for every choice of x and y there is a walk (path) from x to y ; the *diameter* D_G of a strongly connected graph is the maximum length of a shortest path between two nodes.

We shall occasionally deal with (arc-)coloured graphs: a *coloured graph* (with set of colours C) is a graph endowed with a colouring function $\gamma : A_G \rightarrow C$. For symmetric graphs, we require that there is an involution $(\bar{\quad}) : C \rightarrow C$ such that $\gamma(\bar{a}) = \overline{\gamma(a)}$. A (coloured) graph is *separated* iff it has no parallel arcs (with the same colour). The name originates from the fact that such graphs are separated for the double negation topology in the topos of (coloured) graphs—see [34].

A *graph morphism* $\xi : G \rightarrow H$ is given by a pair of functions $\xi_N : N_G \rightarrow N_H$ and $\xi_A : A_G \rightarrow A_H$ commuting with the source and target maps, that is, $s_H \circ \xi_A = \xi_N \circ s_G$ and $t_H \circ \xi_A = \xi_N \circ t_G$ (again, we shall drop the subscripts whenever no confusion is possible). In other words, a morphism maps nodes to nodes and arcs to arcs in such a way to preserve the incidence relation. (In the case of coloured graphs, we require ξ_A to commute with the colouring function.) A morphism between symmetric graphs is *symmetric* iff it commutes with the symmetries. A morphism is *epimorphic* (or an *epimorphism*) iff ξ_N and ξ_A are both surjective.

An *in-tree* is a graph with a selected node r , the root, and such that every other node has exactly one directed path to the root; if t is a node of an in-tree, we sometimes use $t \rightarrow r$ for denoting the unique path from t to the root. If T is an in-tree, we write $h(T)$ for its *height* (the length of the longest path). Finally, 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. A (*symmetric*) *tree* is a (symmetric) graph with a selected node, the root, such that there is exactly one nonstuttering walk (path) from any node to the root: the notions of height and truncation carry on to this case. Unless otherwise stated, morphisms between trees are required to preserve the root.

2.2 Fibrations and coverings

The central concept we are going to deal with is that of *graph fibration*, a particular kind of graph morphism induced by the notion of fibration between categories (see Section 6).

Definition 2.1 A *fibration* between graphs G and B is a morphism $\varphi : G \rightarrow B$ such that for each arc $a \in A_B$ and for each node $x \in N_G$ satisfying $\varphi(x) = t(a)$ there is a unique arc $\tilde{a}^x \in A_G$ (called the *lifting of a at x*) such that $\varphi(\tilde{a}^x) = a$ and $t(\tilde{a}^x) = x$.

We inherit some topological terminology. If $\varphi : G \rightarrow B$ 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 $x \in N_B$ is the set of nodes of G that are mapped to x , and shall be denoted by $\varphi^{-1}(x)$. A fibre is *trivial* if it is a singleton, that is, if $|\varphi^{-1}(x)| = 1$. A fibration is *nontrivial* if at least one fibre is nontrivial, *trivial* otherwise; it is *proper* if all fibres are nontrivial.

There is a very intuitive characterization of fibrations based on the concept of local in-isomorphism. An equivalence relation \simeq between the nodes of a graph G satisfies the *local in-isomorphism property* if the following holds:

Local In-Isomorphism Property: If $x \simeq y$ there exists a (colour-preserving, if G is coloured) bijection $\psi : G(-, x) \rightarrow G(-, y)$ such that $s(a) \simeq s(\psi(a))$, for all $a \in G(-, x)$.

The following proposition shows that fibrations and epimorphisms whose fibres satisfy the previous property are naturally equivalent:

Theorem 2.1 Let G be a graph. Then:

1. if $\varphi : G \rightarrow B$ is a fibration, then the equivalence relation on the nodes of G whose equivalence classes are the nonempty fibres of φ satisfies the local in-isomorphism property;
2. if \simeq is a relation satisfying the local in-isomorphism property, then there exists a graph B and an epimorphic fibration $\varphi : G \rightarrow B$ whose fibres are the equivalence classes of \simeq .

Proof. 1. For each $x, y \in N_G$ such that $x \simeq y$ (i.e., $\varphi(x) = \varphi(y)$) define $\psi : G(-, x) \rightarrow G(-, y)$ by letting $\psi(a) = \tilde{\varphi(a)}^y$. Then we obtain $\varphi(s(\psi(a))) = \varphi(s(\tilde{\varphi(a)}^y)) = s(\varphi(\tilde{\varphi(a)}^y)) = s(\varphi(a)) = \varphi(s(a))$, hence $s(\psi(a)) \simeq s(a)$, as required.

2. Let the bijections $\psi_{x,y}$, whose existence is guaranteed by the local in-isomorphism property, be fixed for every x, y such that $x \simeq y$. Define B as having set of nodes N_G/\simeq , fix a choice of representatives for \simeq , and set

$$B([x], [y]) = \sum_{z \in [x]} G(z, y),$$

where x and y run through the representatives. The definition does not depend on the choice of the representatives because of the local in-isomorphism property.

The map φ is defined on the nodes as $\varphi(x) = [x]$, and on the arcs as follows: let a be an arc of G and $t(a) \simeq y$, where y is a representative; then,

$$\varphi(a) = \psi_{t(a), y}(a).$$

By using the local in-isomorphism property it is now straightforward to show that φ is an epimorphic fibration. ■

Another possible, more geometric way of interpreting the definition of fibration is that given a node x of B and path π terminating at x , for each node y of G in the fibre of x there is a unique path terminating at y that is mapped to π by the fibration; this path is called the *lifting of π at y* , and it is denoted by $\tilde{\pi}^y$. In Figure 1, fibres are represented by dotted ovals (not all nodes of a fibre are shown, though), and we indicate how a path can be lifted at two different points of a fibre. Observe that loops are not necessarily lifted to loops.

It is worth noticing that some simple path-lifting techniques give the following proposition, whose proof is remarkably similar to its topological counterpart.




Figure 1: Different liftings of a path.

Proposition 2.1 A fibration with strongly connected base and nonempty total graph is an epimorphism.

Proof. Let $\varphi : G \rightarrow B$ be a fibration with B strongly connected. Let x be a node of B , π be a path from x to a node y that is in the image of φ (at least one such node exists, being G nonempty), and z be an element of the fibre of y . Then the lifting $\tilde{\pi}^z$ starts from a node in the fibre of x . Surjectivity on arcs follows directly by lifting. ■

A covering projection is a special kind of fibration, where each arc can also be lifted uniquely from its tail; this fact can be seen as the categorical dual of the local in-isomorphism property—more formally:

Definition 2.2 An *opfibration* between graphs G and B is a morphism $\varphi : G \rightarrow B$ such that for every arc $a \in A_B$ and every node $x \in N_G$ satisfying $\varphi(x) = s(a)$, there is a unique arc ${}^x\tilde{a} \in A_G$ (called the *oplifting of a at x*) such that $\varphi({}^x\tilde{a}) = a$ and $s({}^x\tilde{a}) = x$. A *covering projection* is a fibration that is also an opfibration.

If a covering projection $\varphi : G \rightarrow B$ exists, G is said to be a *covering* of B . In the case of coverings, we have a *local isomorphism property* that gives a bijective correspondence between the whole (disjoint) neighbourhoods of two nodes in the same fibre (but note that Theorem 2.1 does *not* generalize—the map one obtains is a fibration, but not in general a covering projection). Covering projections enjoy the following property:

Proposition 2.2 A covering projection $\varphi : G \rightarrow B$ with connected base and nonempty covering is an epimorphism; moreover, the cardinality of all fibres is the same.

Proof. The first part follows as in the proof of Proposition 2.1, using walks instead of paths. Moreover, for every pair of nodes x and y of B , the liftings of a walk from y to x at every node in the fibre of x induce an injection $\varphi^{-1}(x) \rightarrow \varphi^{-1}(y)$, so $|\varphi^{-1}(x)| \leq |\varphi^{-1}(y)|$. ■

The third, and last kind of map we study is strictly related to coverings of undirected graphs:




Figure 2: An eight-cycle covering a four-cycle.

Definition 2.3 A covering projection $\varphi : G \rightarrow B$ between two symmetric graphs is a *symmetric covering projection* if and only if it commutes with the symmetries, that is, for all $a \in A_G$ we have $\varphi(\bar{a}) = \overline{\varphi(a)}$.

An analogous definition for fibrations would lead to the same class of maps, since, as it is easy to show, any symmetric fibration is a covering projection (we shall give a very general categorical proof of this fact in Section 6). It is important to note that “classical” coverings [14] between loopless separated undirected graphs are symmetric coverings in the above sense (assuming undirected graphs are represented as symmetric digraphs), and *viceversa*. However, the situation gets subtler in the case loops are present. This point will be fully discussed in Section 5.

2.3 Groups, fibrations and automorphisms

There is an important relation between fibrations and actions over G . A *left action* ${}_{\Gamma}G : \Gamma \times G \rightarrow G$ of a group Γ on a graph G is a group homomorphism $\Gamma \rightarrow \mathbf{Aut}(G)$. The action is said to be *faithful* if the homomorphism is injective; all actions in this paper are such. We denote the action by left juxtaposition, and ambiguously leave the action name partially unspecified.

The action ${}_{\Gamma}G$ induces an equivalence relation both on the nodes and on the arcs of G , whose classes (the *orbits*) are denoted by $\Gamma(x)$ or $\Gamma(a)$. Note that if x, y are two nodes of G belonging to the same orbit (i.e., $gx = y$ for some $g \in \Gamma$), then the action of g gives a bijection between $G(-, x)$ and $G(-, y)$ that fulfills the requirements of the local in-isomorphism property; thus, by Theorem 2.1, ${}_{\Gamma}G$ induces a fibration $\varphi : G \rightarrow B_{\Gamma}$, where B_{Γ} is a graph having as node set the set of orbits of ${}_{\Gamma}G$ and as many arcs from orbit $\Gamma(x)$ to orbit $\Gamma(y)$ as the arcs coming into an element of $\Gamma(y)$ from all elements of $\Gamma(x)$. We say that φ is *associated with* ${}_{\Gamma}G$. Not all fibrations are associated with an action: a cubic (i.e., 3-regular) graph with trivial automorphism group is fibred over a graph with one node and three loops, but it has no nontrivial associated fibrations.

Note that φ is in general not unique, as it depends, for every x and y , on the element of Γ that is chosen to induce the local in-isomorphism between $G(-, x)$ and $G(-, y)$. Moreover, between G and B_{Γ} there could exist other fibrations that cannot be constructed in this way. Even more is true: some, but not all, of the fibrations associated to an action could happen to be covering projections, as one can easily see by considering the fibrations associated with the action of the automorphism group of a bidirectional 3-cycle.

There is, of course, a more standard object associated to ${}_{\Gamma}G$, viz., the *quotient graph* G/Γ , whose nodes and arcs are the orbits of nodes and arcs of G under the action ${}_{\Gamma}G$.

An action is said to be *free* or *semiregular* iff for all nodes x, y there is at most one $g \in \Gamma$ such that $gx = y$ (note that this fact implies the same for arcs). Equivalently, one can require that no

element of $\Gamma \setminus \{1\}$ has fixpoints. The quotient projection $\Gamma : G \rightarrow G/\Gamma$ is a covering projection under the hypothesis that ΓG is *free*. Conversely, if G is connected and the map Γ is a covering projection then ΓG is free.

The action ΓG induces an epimorphism $\alpha : B_\Gamma \rightarrow G/\Gamma$, which is the identity on the nodes and maps an arc a to $\Gamma(a)$ (recall that a is also an arc of G). The following commutative diagram shows the relation between the aforementioned maps:

$$\begin{array}{ccc} G & & \\ \varphi \downarrow & \searrow \Gamma & \\ B_\Gamma & \xrightarrow{\alpha} & G/\Gamma \end{array}$$

It is possible to characterize the actions for which α is an isomorphism, as follows:

Proposition 2.3 The map $\alpha : B_\Gamma \rightarrow G/\Gamma$ is an isomorphism iff the action ΓG satisfies the following property: if $g \in \Gamma$ fixes x , then it is the identity on $G(-, x)$ (i.e., it fixes pointwise the arcs coming into x).

Proof. First of all, note that α is always surjective on the arcs, and trivially bijective on the nodes. If ΓG satisfies the abovementioned property, then for every pair of arcs a, b of B_Γ such that $\Gamma(a) = \Gamma(b)$ we have $ga = b$ for some $g \in \Gamma$, and thus, by hypothesis, $a = b$, since they have common target.

On the other hand, if α is an isomorphism and there is a $g \neq 1$ that fixes x but it is not the identity on $G(-, x)$ we have $ga = b$ for some $a, b \in G(-, x)$, so $\Gamma(a) = \Gamma(b)$, contradicting the injectivity of α . ■

As we already remarked, the fibrations associated with an action need not be covering projections, in general. However, this is true if ΓG is free, because in this case by the previous proposition we have $\varphi = \alpha^{-1} \circ \Gamma$, and coverings compose. Note that the freeness hypothesis yields also the uniqueness of φ (there is at most one element of the group that can induce a local isomorphism).

One could wonder whether covering projections are associated to free actions only, but this is false: the automorphism group of a complete graph acts nonfreely, yet an associated fibration is a covering projection. A weaker conjecture could sound as follows: every covering projection associated to an action is also associated to a *free* action (having isomorphic base). This can also be shown false, by the following argument: Petersen's graph (or, more precisely, its representation as a symmetric graph) covers a 3-bouquet by means of a covering projection associated with the action of the automorphism group; nonetheless, no free action can have a single node orbit, for otherwise the graph would be a Cayley graph (see Section 5) by Sabidussi's Theorem.

However, there is a special case in which we can reverse the implication; we call an automorphism of G *node-trivial* iff it acts as the identity on the nodes (such an automorphism may only permute parallel arcs).

Proposition 2.4 Let G be a connected graph. If Γ is cyclic, no $g \in \Gamma \setminus \{1\}$ is node-trivial and φ is a covering projection then ΓG is free.

Proof. Let g be a generator of Γ . Assume by contradiction that the action is not free; then there is a node x of G and a least $k > 0$ such that $g^k x = x$ and $g^k \neq 1$. Thus, the orbit of x contains k nodes. Since φ is a covering projection, *all* its fibres contain k nodes, hence g^k is node-trivial, which fact is absurd. ■

The previous proposition happens to be particularly useful when the action is generated by a single automorphism of G . (Note that for separated graphs the node-triviality hypothesis can be dropped.) Finally, if G is symmetric and ΓG , besides being free, respects the symmetry of G (i.e., $g\bar{a} = \overline{ga}$ for all $g \in \Gamma$ and $a \in A_G$), the quotient graph is endowed with a natural symmetry, and the associated covering projection is symmetric. The reader should note that there are symmetric free actions that do not correspond to any “undirected” free action as defined usually [14]. For instance, $\mathbf{Aut}(K_2)$ acts freely in our sense on K_2 (seen as a symmetric graph), giving as quotient a single semi-edge. Correspondingly, the unique morphism from K_2 to the quotient is a symmetric covering.

The fact that actions on a graph induce divisors was noted by Petersdorf and Sachs [26], and rediscovered by Schwenk [32]. Jonathan Gross and Thomas Tucker [14] study free actions on undirected graphs, and call *regular* a covering projection that is the quotient projection induced by a free action. In this case, G/Γ and B_Γ coincide by Proposition 2.3, so in our terminology a covering projection is regular iff it is associated with a free action.

2.4 Some properties of fibrations between finite graphs

Sometimes, in the finite case, it is possible to derive special properties of a fibration $\varphi : G \rightarrow B$ as consequences of properties of the total graph G (and of some connectedness assumptions on B). This is most useful in applications, and we collect here three results along this line. We say that a graph is coloured *deterministically* iff the restriction of the colouring function to $G(x, -)$ is injective for all nodes x , that is, iff the automaton with transition graph G is deterministic.

Proposition 2.5 If G is a finite deterministically coloured graph, and B is strongly connected, then every colour preserving fibration $\varphi : G \rightarrow B$ is a covering projection.

Proof. Let x and y be any pair of nodes of B . One can easily build an injection from the fibre of y to the one of x by lifting a path connecting x to y at each element of the fibre of y and taking the starting node of the resulting path. This association is necessarily injective, for otherwise two arcs with the same label should exit from a node along the path. This implies $|\varphi^{-1}(y)| \leq |\varphi^{-1}(x)|$ for all x and y , so every fibre has the same cardinality k .

Let now a be an arc from x to y . Then a is lifted k times along the fibre over y , and this k arcs must start from k distinct nodes in the fibre over x (by determinism, no two arcs with the same colour can exit from the same node). By pigeonholing this implies that a can also be uniquely oplifted. ■

Proposition 2.6 If G is a finite symmetric deterministically coloured graph, and B is strongly connected, then for every colour preserving fibration $\varphi : G \rightarrow B$ we have that B is endowed with a symmetry, and φ is a symmetric covering projection.

Proof. By Proposition 2.5, φ is certainly a covering. We have to show that B is a symmetric coloured graph, and that φ commutes with the symmetries of G and B .

Consider an arc a of B going from x to y . Let z be an element of the fibre over y , and \tilde{a}^z the corresponding lifting of a . Then we define $\bar{a} = \varphi(\overline{\tilde{a}^z})$; in other words, we lift a , we take the symmetric in G , and we map it with φ in B ; note that this process is not dependent on the choice of z , for otherwise G would not be deterministically coloured.

The symmetry we have defined on B is an involution commuting with the symmetry on the colours, as for any arc a of B

$$\gamma(\bar{a}) = \gamma(\varphi(\overline{\tilde{a}^z})) = \gamma(\overline{\tilde{a}^z}) = \overline{\gamma(\tilde{a}^z)} = \overline{\gamma(\varphi(\tilde{a}^z))} = \overline{\gamma(a)}.$$

The fact that φ is a symmetric covering is now trivial, since by definition

$$\varphi(\bar{a}) = \varphi(\widetilde{\varphi(a)^{t(a)}}) = \overline{\varphi(a)}. \blacksquare$$

Proposition 2.7 If G is a finite symmetric graph and B is connected, then given fibrations $\varphi, \psi : G \rightarrow B$ we have $|\varphi^{-1}(x)| = |\psi^{-1}(x)|$ for all nodes x of B .

Note that the fibrations are not assumed (and are not necessarily) symmetric. The proof of the previous proposition is based on the following (trivial) lemma:

Lemma 2.1 Let $\varphi : G \rightarrow B$ be a fibration, where G is a symmetric graph. For all nodes x and y of B let $d_{xy} = |\{a \in A_B \mid s(a) = x \text{ and } t(a) = y\}|$. Then $|\varphi^{-1}(x)|d_{xy} = |\varphi^{-1}(y)|d_{yx}$ holds for all $x, y \in N_B$.

Proof (of Proposition 2.7). Let k be the number of nodes of B , and m_1, m_2, \dots, m_k the cardinality of the fibres of $\varphi : G \rightarrow B$. Since B has at least $k - 1$ distinct (unordered) pairs of connected nodes (by connection), the previous lemma gives us at least $k - 1$ independent homogeneous linear constraint on the m_i 's. Then the equation $m_1 + m_2 + \dots + m_k = |N_G|$ forces the system to have at most one solution. \blacksquare

2.5 An application

Our study of graph fibrations was inspired by a problem in distributed computing. Consider a finite strongly connected graph G , whose nodes we shall call *processors*. Each processor has an internal state belonging to a set X , and unlimited computational power. During a step of computation each processor changes its state depending on its own state and on the states of its in-neighbours, that is, the arcs represent unidirectional links along which a processor transmits its state (the change of state may also depend on the colours of the arcs; more precisely, the transition function depends on the multiset of pairs $\langle c, x \rangle$, where c is the colour of an arc coming into a processor and x the state of the processor at the other end). All processors change state at the same time.

One of the main problems of such distributed networks is to establish which configurations of states can be reached when all processors start from the same state and run the same algorithm, or, as usually stated, when the network is *anonymous* (or *uniform*). The main point to be noted here is that, under such constraints, *the existence of a fibration $G \rightarrow B$ forces all processors in the same fibre to remain always in the same state*.

This fact is of particular importance for a number of problems, for instance, the paradigmatic *election problem*, which asks for an algorithm leaving the network in the following state: exactly one processor in a state b (elected), and all other processors in a state a (non-elected). It is clear that if a proper fibration $G \rightarrow B$ exists, no algorithm will ever be able to solve this problem on G . The study of universal total graphs, carried out in the following section, has made it possible to turn this condition into a necessary and sufficient one.

The study of (symmetric) coverings is fundamental for the classification of graphs that admit election algorithms *under certain assumptions on the communication primitives*. Assume for instance that we have a deterministically coloured graph, but we want to write an election algorithm working independently of any particular colouring. In a real-world model this corresponds, using a simple emulation algorithm, to the assumption that each processor is able to distinguish its outgoing links, that is, it is able to send a specific, different message along different links.

We know that, for each particular colouring, election is possible iff the resulting coloured graph has no proper fibration. On the other hand, we have just shown that such fibrations will really be coverings; thus, the existence of a colouring of the graph inducing a proper fibration shows that the graph is a proper covering. But, conversely, a graph that is a proper covering has a deterministic colouring inducing a proper fibration (it can be obtained by colouring deterministically the projection base and lifting the colours—note that to do this a divisor would not suffice: we actually need a graph morphism), whence we conclude that *the networks admitting an election algorithm with the assumption of distinguished outcoming links are exactly those whose underlying graph is covering prime* (i.e., it does not cover nontrivially another graph—see Section 4).

An analogous reasoning shows that, under the assumption of fully bidirectional links, the networks admitting an election algorithm are exactly those whose underlying graph is *symmetric-covering prime*. For more details, see [5].

It is also of interest to consider the *central daemon* model, in which exactly one processor is activated at a time. The order of activation is not known, and an algorithm solving election in this model must work no matter which order is actually selected. We shall not go into detail here, but it is possible to prove that the existence of a fibration $G \rightarrow B$ such that the strong components of the subgraphs induced by fibres are singletons forces all processors in the same fibre always to remain in the same state.

3 Universal fibrations and coverings

In this section we prove the existence of certain trees fibred over a graph G that give the “largest” possible fibration, in a sense that will be made precise by the following theorems.

3.1 Universal total graphs

Theorem 3.1 Let T be an in-tree with root r , and let $\nu : T \rightarrow G$ be a fibration. Then, for each fibration $\varphi : H \rightarrow G$, there exist exactly $|\varphi^{-1}(\nu(r))|$ fibrations $\psi : T \rightarrow H$ such that $\varphi \circ \psi = \nu$; more precisely, the fibration ψ is uniquely determined by the choice of $\psi(r)$ in the set $\varphi^{-1}(\nu(r))$.

Proof. For each $y \in \varphi^{-1}(\nu(r))$ we shall define a map ψ_y as follows: a node t of T is mapped to $s(\widetilde{\nu(t \rightarrow r)}^y)$ (the source of the path obtained by lifting $\nu(t \rightarrow r)$ to y); the map on arcs is defined in the obvious way, that is, $\psi_y(a) = \widetilde{\nu(a)}^{\psi_y(t(a))}$. Note that $\psi_y(r) = s(\widetilde{\nu(r \rightarrow r)}^y) = y$, so the maps ψ_y are all distinct; moreover ψ_y is a fibration. Finally

$$\varphi(\psi_y(t)) = \varphi(s(\widetilde{\nu(t \rightarrow r)}^y)) = s(\varphi(\widetilde{\nu(t \rightarrow r)}^y)) = s(\nu(t \rightarrow r)) = \nu(s(t \rightarrow r)) = \nu(t),$$

as required. Now, let $\psi : T \rightarrow H$ be any fibration such that $\varphi \circ \psi = \nu$, and take $y = \psi(r)$; necessarily $\nu(r) = \varphi(\psi(r)) = \varphi(y)$, so $y \in \varphi^{-1}(\nu(r))$, and one immediately verifies that $\psi = \psi_y$. ■

Thus, every fibration of an in-tree to a graph G is universal, in the sense that essentially every other fibration with base G factors it (a categorical characterization of such fibrations in terms of adjoint functors will be given in Section 6). Note that the tree T of Theorem 3.1 is unique up to the choice of $\nu(r)$:

Corollary 3.1 Let T, T' be two in-trees with roots r, r' , and let $\nu : T \rightarrow G$ and $\nu' : T' \rightarrow G$ be two fibrations. If $\nu(r) = \nu'(r')$ then $T \cong T'$.

Proof. Using Theorem 3.1, we obtain a fibration $\psi : T \rightarrow T'$ such that $\psi(r) = r'$. But such fibration is necessarily an isomorphism, since T and T' are in-trees. ■

We shall now prove that such “universal” fibration exists:

Theorem 3.2 For every node x of a graph G there is an in-tree \tilde{G}^x with root r , and a fibration $v_G^x : \tilde{G}^x \rightarrow G$, such that $v_G^x(r) = x$; we call v_G^x the *universal fibration of G at x* , and \tilde{G}^x the *universal total graph of G at x* .

Proof. We define the in-tree \tilde{G}^x as follows:

- the nodes of \tilde{G}^x are the finite paths of G ending in x ;
- there is an arc from the node π to the node π' iff $\pi = a\pi'$ for some arc a (if G is coloured, then the arc gets the same colour as a).

We then define the graph morphism v_G^x from \tilde{G}^x to G by mapping each node π of \tilde{G}^x (i.e., each path of G ending in x) to its starting node, and each arc of \tilde{G}^x to the corresponding arc of G . It is immediate to check that v_G^x is a fibration. ■

Observe that, by the universal property, for every fibration $\varphi : H \rightarrow G$ and for every node $y \in \varphi^{-1}(x)$ there is a unique isomorphism $\iota : \tilde{G}^x \rightarrow \tilde{H}^y$ such that the following diagram commutes:

$$\begin{array}{ccc}
 \tilde{G}^x & \xrightarrow{\iota} & \tilde{H}^y \\
 \downarrow v_G^x & & \downarrow v_H^y \\
 & & H \\
 & \swarrow \varphi & \\
 G & &
 \end{array}$$

3.2 Universal coverings

Similar properties relate covering projections and trees:

Theorem 3.3 Let T be a (symmetric) tree with root r , G be a (symmetric) graph and $v : T \rightarrow G$ be a (symmetric) covering projection. Then, for each (symmetric) covering projection $\varphi : H \rightarrow G$, there exist exactly $|\varphi^{-1}(v(r))|$ (symmetric) coverings projections $\psi : T \rightarrow H$ such that $\psi \circ \varphi = v$; more precisely, ψ is uniquely determined by the choice of $\psi(r)$ in the set $\varphi^{-1}(v(r))$.

The proof is similar to that of Theorem 3.1, and thus omitted. Correspondingly, we have a notion of (*symmetric*) *universal covering* of G at x , obtained by replacing the paths into x by the nonstuttering walks (paths) into x :

Theorem 3.4 For every node x of a (symmetric) graph G there is a (symmetric) tree \overline{G}^x , with root r , and a (symmetric) covering $\pi_G^x : \overline{G}^x \rightarrow G$, such that $\pi_G^x(r) = x$.

Proof. We just discuss the symmetric case, the other one being similar. Define \overline{G}^x as $\mathbf{Sym}(T)$, where T is the subtree of \tilde{G}^x induced by those paths that are not symmetrically stuttering; \overline{G}^x is defined by extending v_G^x in the natural way (as $\mathbf{Sym}(-)$ is an adjoint functor—see Section 6). Thus, for example, there will be an arc from the node $a\pi$ to the node π , and another arc going in the opposite direction, added by symmetrization. The former will be mapped by π_G^x to a , and the latter to \bar{a} .

Clearly π_G^x is symmetric, so we just have to show that it is actually a fibration; let now a be an arc of G , and π a path from $t(a)$ to x . If π does not start with \bar{a} , then a can be lifted to the only arc going from $a\pi$ to π ; conversely, if $\pi = \bar{a}\pi'$, then a can be lifted to the only arc going from π' to π . ■

Note that if one looks at \overline{G}^x as a (symmetric) graph, then it depends only on the connected component in which x lies, that is, \overline{G}^x is isomorphic to \overline{G}^y , for every choice of x and y in the same component. On the other hand, if we look at \overline{G}^x as a graph with a selected node (the root), then different nodes will usually possess different universal (symmetric) coverings: we shall use the term “rooted universal (symmetric) covering” in this case. A purely combinatorial proof of the invariance of (symmetric) universal coverings along connected components is rather cumbersome; however, this fact will be an immediate consequence of a categorical representation theorem given in Section 6.

The construction of the symmetric universal covering corresponds essentially to the standard construction of a universal covering of undirected graphs from topological graph theory; however, it also includes naturally the case of loops (fixed or otherwise), upon which there seems to be little agreement.

Indeed, our definition solves the rather subtle issues determined by the presence of loops: for instance, a symmetric graph with one node and two loops has the bidirectional line as universal symmetric covering, but the universal symmetric covering projection is different depending on whether the symmetry is the identity or not; moreover, when only one loop is present the universal symmetric covering reduces to a single bidirectional segment, which accounts for the “loops counted once vs. loops counted twice” dilemma in the definitions found in the literature.

3.3 Nodes with the same universal total graph

We shall be interested in identifying the nodes of a graph sharing the same universal total graph (or covering), possibly in an effective way. We start with the following “compactness lemma for trees”:

Lemma 3.1 Let T and U be locally finite in-trees (trees, symmetric trees). Then, $T \cong U$ iff $T \upharpoonright k \cong U \upharpoonright k$ for all $k \in \mathbf{N}$.

Proof. Let $\alpha_k : T \upharpoonright k \rightarrow U \upharpoonright k$, for $k \in \mathbf{N}$, be the isomorphisms of the hypothesis. Note that given an isomorphism $\beta : T \upharpoonright j \rightarrow U \upharpoonright j$ such that the set $I = \{i \mid \alpha_i \text{ extends } \beta\}$ is infinite, it is always possible to extend β to an isomorphism $\gamma : T \upharpoonright (j+1) \rightarrow U \upharpoonright (j+1)$, leaving the set $\{i \mid \alpha_i \text{ extends } \gamma\}$ infinite, as the equivalence relation defined on I by $i \simeq i'$ iff the restriction of α_i and $\alpha_{i'}$ to $T \upharpoonright (j+1)$ coincide has finite index. This allows to define by recursion a sequence β_0, β_1, \dots of isomorphisms $\beta_k : T \upharpoonright k \rightarrow U \upharpoonright k$ such that β_{k+1} extends β_k , inducing an isomorphism $T \rightarrow U$. ■

Generalizing in a natural way the classical definition [3] to our setting, we say that a *highly recursive graph* is a locally finite graph G in which N_G and A_G are recursive subsets of \mathbf{N} , the function $G(-, -)$ from $N_G \times N_G$ to the finite subsets of A_G is recursive, and there is a recursive function ν from N_G to the finite subsets of N_G such that $y \in \nu(x)$ iff $G(x, y) \cup G(y, x)$ is nonempty (i.e., the neighbourhood of each node is recursively computable). For symmetric graphs, we also require that the symmetry is a recursive function. The previous lemma, which is of course true also of (symmetric) trees, has an immediate consequence:

Theorem 3.5 Given a highly recursive graph G , the question whether two nodes have different universal total graphs (rooted coverings, rooted symmetric coverings) is semi-decidable, but not decidable.

Proof. Given nodes x and y , using iteratively the function $G(-, -)$ it is possible to build $\tilde{G}^x \upharpoonright k$ and $\tilde{G}^y \upharpoonright k$, and thus semi-decide whether $\tilde{G}^x \upharpoonright k \cong \tilde{G}^y \upharpoonright k$ for some k , so by Lemma 3.1 we can semi-decide whether $\tilde{G}^x \cong \tilde{G}^y$. On the hand, by coding the configurations of a universal Turing machine into natural numbers and putting an arc into $G(x, y)$ whenever x is the next state after y , we obtain a highly recursive graph in which \tilde{G}^x is an infinite path iff the universal Turing machine does not stop starting from the configuration coded by x . By choosing a fixed node z coding a configuration that is known to be nonterminating, we obtain a reduction of the halting problem to the different total graph problem. The proof in the case of (symmetric) coverings is analogous. ■

Consider now the equivalence relations \simeq_k on the nodes of an arbitrary graph G defined by $x \simeq_k y$ iff $\tilde{G}^x \upharpoonright k \cong \tilde{G}^y \upharpoonright k$.

Lemma 3.2 If $\simeq_k = \simeq_{k+1}$ for some $k \in \mathbf{N}$, then $\simeq_k = \simeq_{k+j}$ for all $j \in \mathbf{N}$.

Proof. Let $x \simeq_{k+1} y$, and $\alpha : \tilde{G}^x \upharpoonright (k+1) \cong \tilde{G}^y \upharpoonright (k+1)$ be the corresponding isomorphism. For every arc a coming into the root of \tilde{G}^x , we have $\tilde{G}^{v_G^x(s(a))} \upharpoonright k \cong \tilde{G}^{v_G^y(s(\alpha(a)))} \upharpoonright k$, that is, $v_G^x(s(a)) \simeq_k v_G^y(s(\alpha(a)))$, but this implies the same at depth $k+1$, so there is an isomorphism $\tilde{G}^{v_G^x(s(a))} \upharpoonright (k+1) \rightarrow \tilde{G}^{v_G^y(s(\alpha(a)))} \upharpoonright (k+1)$. Combining these isomorphisms for all a , we obtain an isomorphism $\tilde{G}^x \upharpoonright (k+2) \cong \tilde{G}^y \upharpoonright (k+2)$, so $x \simeq_{k+2} y$. The result follows by induction. ■

An analogous statement holds, of course, for universal (symmetric) coverings, by redefining suitably \simeq_k . This fact allows one to decide effectively universal total graph isomorphism when a graph is finite, and indeed a result of Nancy Norris [23] could be restated as follows in our terminology:

Theorem 3.6 If G has n nodes, for all nodes x, y , $\tilde{G}^x \cong \tilde{G}^y$ iff $\tilde{G}^x \upharpoonright (n-1) \cong \tilde{G}^y \upharpoonright (n-1)$, that is, iff there is an isomorphism between the first $n-1$ levels of the two trees. The same holds for rooted universal coverings.

Here, we extend the previous theorem to universal symmetric coverings; we also provide a much shorter proof.

Theorem 3.7 Given a finite graph G with n nodes, two nodes x and y have the same universal total graph (rooted covering, rooted symmetric covering) iff $x \simeq_{n-1} y$. Thus, the question whether two nodes have the same universal total graph (rooted covering, rooted symmetric covering) is decidable.

Proof. By compactness, $\tilde{G}^x \cong \tilde{G}^y$ iff $x \simeq_k y$ for all $k \in \mathbf{N}$. But since \simeq_{k+1} refines \simeq_k , certainly $\simeq_{n-1} = \simeq_n$ by the previous lemma and by finiteness of G . ■

The bound given by Theorem 3.6 is tight, as remarked in [23], by the example shown in Figure 3. The very same example shows that even in our case the bound remains tight—in fact, it shows that this is true even if we require the graph to be symmetric and deterministically coloured: the two leftmost nodes share the first $n-2$ levels of their universal total graphs (rooted [symmetric] coverings), but not the first $n-1$.

We get back to the example given in Section 2.5. The reader will have probably guessed at this point that in an anonymous distributed system *processors with the same universal total graph always remain in the same state*. This was already noted (for the undirected case) in the seminal paper by Angluin [2], where she showed that processors with the same universal *covering* always remained in the same state (the network




Figure 3: A graph showing the tightness of Theorem 3.7.

model used there is however much stronger, requiring bidirectionality and determinism of the arc colouring). This condition, which was just shown to be sufficient, turned out also to be essentially necessary later, in the work of Yamashita and Kameda [37], who considered same model.

The constructive proof of Theorem 3.7 can be easily modified to show that processors sharing the first k levels of their universal total graphs have the same state during the first k steps of the computation, whichever algorithm and initial state one chooses. This implies that no algorithm will ever be able to drive processors having the same universal total graph into distinct states. This is true independently of any other structural property of the graph, and thus applies to a very wide variety of models (for instance, to *wireless* networks, in which it is impossible to distinguish which link provided which message).

On the other hand, there is an algorithm that allows each processor to compute (a finite number of levels of) its own universal total graph. The algorithm is fairly obvious, and consists in reading (at step k) from all neighbours their universal total graph truncated at depth k ; this makes each processor capable of building its universal total graph truncated at depth $k + 1$, and so on (for more details, see [6]). Once enough levels are known, each processor knows which equivalence class (of isomorphism of total graphs) it lies in; then, for instance, the processors in classes that are singletons can initiate a standard election algorithm (e.g., by lexicographic ordering of the universal total graph).

The situation is much more complicated in the central daemon case, or if we consider a class of network instead of a single network. The theory of minimum bases and minimal fibrations, developed in the next section, approaches exactly these problems.

4 Minimal fibrations

It is worth noticing that every fibration of a graph “smashes together” some nodes that possess the same universal total graph:

Proposition 4.1 If $\varphi : G \rightarrow B$ is a fibration, and $\varphi(x) = \varphi(y)$, then $\tilde{G}^x \cong \tilde{G}^y$.

Proof. By the universal property, we have $\tilde{B}^{\varphi(x)} \cong \tilde{G}^x$ and also $\tilde{B}^{\varphi(y)} \cong \tilde{G}^y$, from which the conclusion follows immediately. ■

It is natural to ask whether it is possible to take this process to extremes and identify any two nodes having the same universal total graph. This question will be answered shortly, after some definitions are introduced. A graph G is *node-rigid* iff every automorphism of G is node-trivial; it is *rigid* iff its automorphism group is trivial.

Definition 4.1 A graph G is *fibration prime* iff it cannot be fibred nontrivially, that is, every epimorphic fibration $G \rightarrow B$ is an isomorphism.

In previous papers [5, 6, 7] fibration-prime graphs have been called *trivial bundles*. The present change of terminology was dictated by the desire of avoiding confusion with the current topological custom. Moreover, it pays a tribute to divisor theory—graphs without rear divisors are exactly fibration-prime graphs.

Proposition 4.2 A fibration-prime graph is node-rigid. A separated fibration-prime graph is rigid.

Proof. Let G be fibration prime, and suppose α is a non–node-trivial automorphism of G . Then every fibration associated with the action of the subgroup generated by α is nontrivial—a contradiction. Finally, note that for separated graphs an automorphism is node-trivial iff it is the identity. ■

Theorem 4.1 Let G be a graph. Then there exists a graph B such that G is epimorphically fibred over B , and the universal total graphs of B are pairwise nonisomorphic.

Proof. Define $x \simeq y$ iff $\tilde{G}^x \cong \tilde{G}^y$. Then \simeq enjoys the local in-isomorphism property, and the claim follows by Theorem 2.1. ■

This leads to the useful

Corollary 4.1 A graph is fibration prime iff its universal total graphs are pairwise nonisomorphic.

Proof. If G has a pair of isomorphic universal total graphs, by Theorem 4.1 we can fibre it non-trivially. The other direction is immediate by Proposition 4.1. ■

The important property of fibration-prime graphs we shall need is given by the following

Theorem 4.2 Let B and C be fibration prime, and suppose they have the same (set of) universal total graphs. Then $B \cong C$, and the node component of such isomorphisms is unique.

Proof. Since by Corollary 4.1 no two nodes of B (C , respectively) have isomorphic universal total graphs, there is a unique bijection $\varphi : N_B \rightarrow N_C$ such that $\tilde{B}^x \cong \tilde{C}^{\varphi(x)}$ for all nodes x of B . Consider now an arc a of B with target x ; the above isomorphism associates to \tilde{a}^r (r is the root of \tilde{B}^x , and the lifting is along v_C^x) a unique arc b of $\tilde{C}^{\varphi(x)}$, and we define $\varphi(a) = v_C^{\varphi(x)}(b)$. Note that the source of \tilde{a}^r induces a subtree of \tilde{B}^x that is isomorphic to the subtree induced by the source of b ; by the abovementioned uniqueness property, this fact ensures that the source of $\varphi(a)$ is the image through φ of the source of a , and because of the local in-isomorphism property φ is an isomorphism. Finally, note that the existence of two isomorphisms with a different node component between B and C would imply the existence of a non–node-trivial automorphism of B (and C), which is impossible by Proposition 4.2. ■

The above theorems suggest to investigate fibrations whose base is fibration prime:

Definition 4.2 A fibration $\mu : G \rightarrow B$ is *minimal* iff it is an epimorphism and B is fibration prime.

Theorem 4.3 If G is minimally fibred over B and C , then there is an isomorphism $\alpha : B \xrightarrow{\sim} C$ and the node component of the two fibrations is the same, modulo composition with (every such) α .

Proof. Let $\varphi : G \rightarrow B$ and $\psi : G \rightarrow C$ be the two fibrations. Clearly B and C have the same universal total graphs as G ; thus, there is an isomorphism $\alpha : B \xrightarrow{\sim} C$. But for each $x \in N_G$

$$\tilde{B}^{\alpha(\varphi(x))} \cong \tilde{B}^{\varphi(x)} \cong \tilde{G}^x \cong \tilde{C}^{\psi(x)},$$

which by primality of C implies $\alpha(\varphi(x)) = \psi(x)$. ■

Thus, all minimal fibrations of a graph G have (up to isomorphism) the same codomain, which is called the *minimum base of G* , and denoted by \hat{G} (an example is given in Figure 4); moreover, they all behave in the same way with respect to the nodes (they can only differ in the way they map




Figure 4: A graph, one of its universal total graphs and its minimum base.




Figure 5: Two nonisomorphic prime coverings with the same universal covering.

arcs of G to arcs of \widehat{G}). Note however that Theorem 4.3 does not extend to (symmetric) coverings, as the following example shows. The two graphs in Figure 5 are (symmetric-) covering prime, that is, they cannot cover nontrivially another graph, for they have a prime number of nodes (see Proposition 2.2). They are also nonisomorphic (the first one gets disconnected by the removal of a node, the second one does not) but nonetheless it is easy to check that they share the same universal (symmetric) covering. In particular, by Leighton’s Theorem ([19]—see Section 6), they share a common finite (symmetric) cover. So we have an example of a graph whose minimal (symmetric) covering bases are not isomorphic.

It is interesting to note that a sufficient condition for a finite graph to be fibration prime is that its characteristic polynomial (i.e., the characteristic polynomial of its adjacency matrix) be irreducible over \mathbf{Z} . This is due to the fact that, as we remarked in the introduction, the existence of a fibration $\varphi : G \rightarrow B$ implies that the characteristic polynomial of B divides the one of G (and the quotient lies in $\mathbf{Z}[x]$). The implication cannot be reversed: the 2-outregular graph with exactly two nodes and one loop has characteristic polynomial $(\lambda - 2)(\lambda + 1)$, although it is fibration prime.

The reader familiar with process algebras and their semantics (see, e.g., [21]) will have certainly noticed that the *unfolding* of a labelled transition system with initial state, that is the *synchronization tree* generated by a labelled graph G with a selected node x , is exactly the graph that is *universally opfibred over G at x* , or, equivalently, the dual of the universal total graph at x of the dual of G . Even more is true: any graph G is *strongly bisimilar* to its minimum base, but if G is not deterministically coloured there can be even smaller graphs that are strongly bisimilar to G , which is probably the reason why *zig-zag morphisms* (see, e.g., [17]) were preferred to fibrations in graph-theoretical formalizations of strong bisimilarity.

4.1 Constructing minimum bases

Theorem 4.1 provides, in the finite case, a constructive procedure for building minimal fibrations; however, a more efficient algorithm can be obtained by using *set-partition* techniques.

Theorem 4.4 Given a finite graph G , there is a set-partition algorithm that computes the minimum base of G and a minimal fibration.

Proof. The algorithm we are going to describe partitions the graph into classes of nodes having the same universal total graph. To do so, it works in $|N_G| - 1$ phases: the partition associated to phase k is the one induced by the equivalence relation \simeq_k described in Section 3.3 (i.e., two nodes are in the same class iff they share the first k levels of their universal total graphs). By Theorem 3.7, after the last phase two nodes are in the same class iff they have the same universal total graphs; this relation enjoys the local in-isomorphism property, and thus induces a fibration, which is minimal by Corollary 4.1.

At phase 0, all nodes are in the same class, because \simeq_0 is the total relation. To build \simeq_{k+1} , just note that $x \simeq_{k+1} y$ iff $x \simeq_k y$ and there is a bijection $\psi : G(-, x) \rightarrow G(-, y)$ such that $s(a) \simeq_k s(b)$ and $t(a) \simeq_k t(b)$, and refine the current partition accordingly. ■

The previous theorem allows one to derive from a graph its minimum base. However, it is possible to build \widehat{G} with much less information. We have already seen (Theorem 3.7) that for finite graphs the isomorphism of universal total graphs needs to be tested only on $n - 1$ levels. Thus, $n + D$ levels (recall that D is the diameter of G) of *any* universal total graph of G contain enough information to rebuild the minimum base, *given the knowledge of n and D* . The following theorem shows that the minimum base can be constructed *even without knowing n and D* . That is, given a sufficiently deep finite truncation of a universal total graph, we can always build the minimum base of its graph without using other information.

Theorem 4.5 Let G be a strongly connected graph with n nodes and diameter D and B a fibration-prime graph with minimum number of nodes satisfying $\widetilde{G}^x \upharpoonright (n + D) \cong \widetilde{B}^y \upharpoonright (n + D)$ for some $x \in N_G$ and $y \in N_B$: then $B \cong \widehat{G}$.

Proof. Note that B has at most n nodes, because the minimum base of G satisfies the hypotheses. We shall build a morphism $\varphi : G \rightarrow B$ by sending a node z of G to the unique node $\varphi(z)$ of B satisfying $\widetilde{G}^z \upharpoonright (n - 1) \cong \widetilde{B}^{\varphi(z)} \upharpoonright (n - 1)$. This node can be found as follows: there is certainly a node $z' \in (v_G^x)^{-1}(z)$ that is at depth D at most. Thus, the subtree under z' in $\widetilde{G}^x \upharpoonright (n + D)$ has height at least $n - 1$. Let $\psi : \widetilde{G}^x \upharpoonright (n + D) \rightarrow \widetilde{B}^y \upharpoonright (n + D)$ be the isomorphism above. Then $\varphi(z) = (v_B^y \circ \psi)(z')$. Note that the choice of z' is irrelevant, by Corollary 4.1.

We now define analogously φ on the arcs, by using the lifting property. Let a be an arc of G . We choose, as before, a $z \in (v_G^x)^{-1}(t(a))$ that is at depth D at most, and consider the lifting \widetilde{a}^z . Then we set $\varphi(a) = (v_B^y \circ \psi)(\widetilde{a}^z)$. Note that this is compatible with our definition on the nodes, because $s(\widetilde{a}^z)$ is at depth $D + 1$ at most, and thus its image through $v_B^y \circ \psi$ must be $\varphi(s(a))$, by Theorem 3.7. 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. ■

The bound given in Theorem 4.5 is tight. Consider the families of graphs $G_{n,D}$ and $H_{n,D}$ (with n nodes and diameter D) depicted in Figure 6 (the only difference between the two families is given by the dotted arc). It is easy to show that $\widetilde{G_{n,D}}^1 \upharpoonright (n + D - 1) \cong \widetilde{H_{n,D}}^1 \upharpoonright (n + D - 1)$, but the graphs are fibration prime. Thus, in general the bound $n + D$ cannot be improved.

The theorems proved in this section allow one to characterize effectively the solvability of the election (and virtually any other computability) problem anonymously. Although the machinery we developed is definitely overkill for a single network, it can be used to provide analogous results for *arbitrary* classes of networks, and also to decide the computability of functions [6] or relations.

The main idea is that if we have a class \mathcal{C} of networks, and we want to know, for instance, whether an election algorithm working for all networks of \mathcal{C} exists, we must study the minimum bases (and related fibrations) of all networks of \mathcal{C} . Essentially, *for each fibration-prime graph B there must be a node x such that for every minimal fibration $\varphi : G \rightarrow B$, with $G \in \mathcal{C}$, the fibre over x is trivial (i.e., $|\varphi^{-1}(x)| = 1$)*. This is a necessary and sufficient condition, and works also for the central daemon case, provided that we restrict the class of fibrations used in the way discussed in Section 2.5.

Sometimes we do not require an election algorithm to terminate—rather, we need a *self-stabilization* property: even if the algorithm is nonterminating, after a finite number of steps the global state of the system is an election state. Theorem 4.5 is a fundamental tool in providing an upper bound for this number of steps in the most general case. Consider an arbitrary infinite class of networks \mathcal{C} such that election is possible in every finite subclass of \mathcal{C} . In this case it is possible to use the algorithm described previously, assuming the existence of a larger number of nodes in the network at each step. No matter how large the network is, after exactly $n + D$ steps the processors will enter an election state, since they will compute correctly the minimum base. These considerations can be pushed further to *every* computable self-stabilizing nonreactive behaviour, as done in [7].

5 Graphs fibred over bouquets

As we discussed in the previous section, sometimes minimal covering bases may not be isomorphic. However, one can still “work backwards” and, given a base graph B that is a prime covering, classify the related covering spaces, that is, the graphs having B as minimal (epimorphic) covering base (one could even work with a set of such bases).

In this section, we attack the simplest case and characterize the graphs defined by the property of being fibred over (or covering) a bouquet (i.e., a graph with exactly one node). It is obvious that the graphs fibred over bouquets are exactly the irregular graphs; if the fibration is required




Figure 6: Graphs with similar universal total graphs.

to be associated with an action, we obtain exactly the node-transitive graphs, and if the action is required to be free we obtain exactly the Cayley graphs, by Sabidussi's theorem [28]. The characterization of (symmetric) coverings of bouquets is however more interesting, and requires sometimes additional hypotheses.

We set up some terminology: a d -factor of a graph G , where $d \in \mathbf{N}$, is a d -regular subgraph containing all nodes of G ; a *symmetric d -factor* is a d -factor closed by symmetry. For a set $I \subseteq \mathbf{N}$, a (symmetric) I -factorization of G is a set of arc-disjoint subgraphs of G , such that each subgraph is a (symmetric) d -factor, for some $d \in I$, and each arc of G belongs to one of the factors (if I is a singleton, we omit curly braces). Note that 1-factors of symmetric graphs are usually called *perfect matchings* (but some care must be taken in interpreting correctly the meaning of semi-edges).

The following theorem highlights the relation between coverings of bouquets and factors:

Theorem 5.1 A graph covers a bouquet iff it is 1-factorable. A symmetric graph covers symmetrically a bouquet iff it is symmetrically $\{1, 2\}$ -factorable in such a way that the 2-factors do not contain semi-edges.

Proof. In the first case, the fibre of a loop of the bouquet induces exactly a 1-factor, by uniqueness of lifting and oplifting. Analogously, in the second case the fibre of a pair of loops exchanged by the symmetry is a symmetric 2-factor without semi-edges, while the fibre of a semi-edge is a symmetric 1-factor.

On the other hand, if a graph is 1-factorable the morphism to a bouquet (with as many loops as there are factors) that sends each factor to a distinct loop is trivially a covering projection. In the symmetric case, we add to the bouquet a pair of loops exchanged by the symmetry for each symmetric 2-factor, and a semi-edge for each 1-factor. For each component of a symmetric 2-factor, we choose arbitrarily to send a selected arc a to one of the two loops: this choice extends uniquely to all other arcs (in the other case the map is unique). ■

Consequently, our classification results will depend on some lemmata about factorizations. We remark that there are no hidden cardinality assumptions; in the rest of this section, we shall use silently the axiom of choice, as we have already done in the proof of Theorem 2.1.

5.1 Factorization lemmata

The first result we need is in fact a well-known matching theorem for finite bipartite undirected graphs, which can be interpreted as an existence theorem for 1-factors in digraphs, and can be extended to the countable, locally finite case [24]. We provide a self-contained proof, which does not depend on matching theory and turns out to be fairly shorter; moreover, we extend the original statement to nonseparated graphs. Given a set of nodes X , we let $G^+(X) = t(G(X, -))$, that is, $G^+(X)$ is the set of nodes that are targets of arcs going out of X . We recall that a *sink* is a node without outgoing arcs (i.e., $G(x, -) = \emptyset$); a subset of nodes is *sink-free* if it does not contain any sink.

Lemma 5.1 Let G be a d -regular graph. Then G has a 1-factor (so it is 1-factorable).

Proof. Given a sink-free finite subset X of N_G let

$$\delta(X) = |X| - |G^+(X)|$$

be the *deficiency* of X (i.e., the difference between $|X|$ and the number of nodes that are targets of arcs going out of X). The deficiency of G , denoted δ , is defined as the supremum among the

deficiencies of all finite sink-free subsets of nodes

$$\delta = \sup\{\delta(X) \mid X \text{ is a finite sink-free subset of } N_G\},$$

and a simple pigeonholing argument shows that regular graphs of finite degree are without deficiency (i.e., $\delta = 0$; note that $\delta(\emptyset) = 0$, so the deficiency of G cannot be negative). Because of the disequation

$$|G^+(X \cup Y)| + |G^+(X \cap Y)| \leq |G^+(X)| + |G^+(Y)|,$$

we have that

$$\delta(X \cup Y) + \delta(X \cap Y) \geq \delta(X) + \delta(Y),$$

so, in particular, in a graph without deficiency the intersection and union of finite sink-free subsets with deficiency 0 have still deficiency 0.

Let now F be a 1-subfactor of G , that is, a subgraph of G such that every node has at most one incoming and at most one outgoing arc. We define $G \searrow F$ as follows: we subtract from G all arcs having the same source or target of an arc in F . Consider now the partially ordered set of 1-subfactors F that leave $G \searrow F$ without deficiency. If we have a totally ordered subset O of such 1-subfactors, $\bigcup O$ is a 1-subfactor, and the graph $G \searrow \bigcup O$ is without deficiency, for these conditions must be false for some element in O if they are false for $\bigcup O$. So there must be a maximal 1-subfactor M leaving G without deficiency, and we are going to show that it is a 1-factor of G (this happens iff $G \searrow M$ has no arcs).

Assume by contradiction that $G \searrow M$ has one or more arcs. We show that we can easily add an arc to M , contradicting maximality. In the rest of the proof, X and Y will always denote finite nonempty sink-free subsets of the graph under consideration.

If all subsets X of $G \searrow M$ have strictly negative deficiency, any arc can be added to M . Indeed, a subset Y of $G \searrow M_a$, where M_a is obtained by adding an arbitrary arc a to M , cannot contain $s(a)$, so

$$(G \searrow M_a)^+(Y) = (G \searrow M)^+(Y) \setminus \{t(a)\},$$

and this implies $\delta(Y) \leq 0$. Otherwise, we consider a minimal subset X of $G \searrow M$ such that $\delta(X) = 0$ in $G \searrow M$, and we choose an arc a going out of some node of X . Given a subset Y of $G \searrow M_a$, we must show that it has nonpositive deficiency. If Y has strictly negative deficiency in $G \searrow M$, the argument goes as in the previous case. Otherwise, Y cannot contain X (because it does not contain $s(a)$), so it is disjoint from X (by minimality of X). But then $t(a) \notin (G \searrow M)^+(Y)$, for otherwise $X \cup Y$ would have positive deficiency in $G \searrow M$, so Y still has deficiency 0 in $G \searrow M_a$. ■

We now extend the previous theorem to regular graphs of infinite degree \mathfrak{c} in a special case: recall that a node y is a *successor* of x iff $G(x, y)$ is nonempty, and a *predecessor* of x if $G(y, x)$ is nonempty; we say that a graph is *well balanced* iff $|t(G(x, -))| = |G(x, -)|$ and $|s(G(-, x))| = |G(-, x)|$ for all nodes x , that is, every node has as many successors as outgoing arcs, and as many predecessors as incoming arcs. On locally finite graph well balancing is equivalent to separatedness, but this is not true when we turn to graphs with infinite local degree (of course, separated implies well balanced).

Lemma 5.2 Let G be a well-balanced \mathfrak{c} -regular graph. Then G is 1-factorable.

Proof. We note that G can be assumed of cardinality \mathfrak{c} (i.e., $|N_G| = |A_G| = \mathfrak{c}$), as by hypothesis each connected component of G must contain at least \mathfrak{c} nodes, and since k -distance neighbours of a node cannot be more than $\mathfrak{c}^k = \mathfrak{c}$, each connected component cannot contain more than $\aleph_0 \mathfrak{c} = \mathfrak{c}$ nodes; clearly, G covers a \mathfrak{c} -bouquet iff all its connected component do.

First of all we prove that every graph H satisfying the hypotheses has a 1-factor including a given arc a . For this purpose, we shall define injective functions f and g on the nodes of H such that $H(x, f(x))$ and $H(g(x), x)$ are nonempty for every node x , and moreover $f(s(a)) = t(a)$, $g(t(a)) = s(a)$. By the Schröder–Bernstein theorem (given injections $f : X \rightarrow Y$ and $g : Y \rightarrow X$, there is a bijection $h : X \rightarrow Y$ such that for all $x \in X$ either $h(x) = f(x)$ or $h(x) = g^{-1}(x)$), we shall then obtain a permutation p on the nodes of H such that $p(s(a)) = t(a)$ and $H(x, p(x))$ is nonempty; as a consequence, we shall be able to build a 1-factor of H including a , by selection of an arbitrary arc from each set $H(x, p(x))$, $x \neq s(a)$.

Let $x_0, x_1, \dots, x_\omega, \dots$ be an ordering of the nodes of H of order type γ , where γ is the least ordinal of cardinality \mathfrak{c} . We build the function f extending by transfinite recursion the definition $f(s(a)) = t(a)$. Given a $\beta < \gamma$ (so $|\beta| < \mathfrak{c}$), and assuming we defined f for all nodes x_α with $\alpha < \beta$, we define $f(x_\beta)$ by choosing arbitrarily the target of an arc coming out of x_β that does not belong to

$$X_\beta = \bigcup_{\alpha < \beta} \{f(x_\alpha)\}.$$

This is always possible, as

$$|X_\beta| = |\{f(x_\alpha) \mid \alpha < \beta\}| = |\{\alpha \mid \alpha < \beta\}| = |\beta| < \mathfrak{c},$$

so there are \mathfrak{c} targets of arcs going out of x_β still left. Of course, the function g can be obtained analogously.

Now, let $a_0, a_1, \dots, a_\omega, \dots$ be an ordering of the arcs of G of order type γ . By transfinite recursion, we define for each $\beta < \gamma$ a subgraph F_β of G that is either empty or a 1-factor. More precisely, given a $\beta < \gamma$, and assuming we defined F_α for all $\alpha < \beta$, we define F_β either as the empty subgraph, if a_β belongs to F_α for some $\alpha < \beta$, or as a 1-factor of the graph obtained deleting from G the arcs in $\bigcup_{\alpha < \beta} F_\alpha$; moreover, we can choose F_β so that it contains a_β , since the deletion of the arcs in $\bigcup_{\alpha < \beta} F_\alpha$ can reduce the cardinality of the successors (predecessors) of a node at most by $|\beta| < \mathfrak{c}$. Clearly the nonempty F_β 's are disjoint, and contain by construction every arc of G . This gives a 1-factorization of G of cardinality \mathfrak{c} and thus, by Theorem 5.1, G covers a \mathfrak{c} -bouquet. ■

Note that well balancing is necessary, as a graph with two nodes x, y and $\mathfrak{c} = |G(x, y)| = |G(y, x)| = |G(x, x)| > |G(y, y)|$ shows.

Symmetric factorizations are more difficult to deal with, and we shall need some additional hypotheses.

Lemma 5.3 Let G be a symmetric locally finite even graph (i.e., a graph such that the indegree of each node is equal to its outdegree, and they are both even) without semi-edges. Then, the arcs of G can be partitioned so that the arcs in each class (and their endpoints) form a symmetric connected 2-regular subgraph of G (i.e., a bidirectional cycle or infinite line).

Proof. First we prove that each graph H satisfying the hypotheses contains a symmetric connected 2-regular subgraph. We can assume without loss of generality that H is separated and loopless, because otherwise we could obtain the subgraph above closing by symmetry two parallel arcs or

a loop. Let now a be an arc going from x_0 to x_1 and, using the even degree assumption, build by recursion a biinfinite sequence of nodes $\dots, x_{-1}, x_0, x_1, x_2, \dots$ such that $x_{i-2} \neq x_i \in H^+(x_{i-1})$. A maximal subsequence without repetitions defines a symmetric connected 2-regular subgraph.

Consider now the set of subpartitions of A_G (i.e., partition of subsets of A_G) whose classes form symmetric connected 2-regular subgraphs of G , partially ordered by inclusion. Every chain in this set has a bound (the union), so there is a maximal element M . If M is not a partition of A_G , we consider the graph having the node set of G but arcs $A_G \setminus \bigcup M$, which satisfy the hypotheses of the theorem, and thus has a symmetric connected 2-regular subgraph, which can be added to M , contradicting its maximality. ■

By using this lemma, one can prove that:

Theorem 5.2 Let G be a symmetric graph without semi-edges. Then:

1. if G is $2m$ -regular, then it is symmetrically 2-factorable;
2. if G is $(2m+1)$ -regular, then it is symmetrically $\{1, 2\}$ -factorable iff it possesses a symmetric 1-factor;
3. if G is well balanced and \mathfrak{c} -regular, then it is symmetrically 2-factorable.

Proof. (1). Consider the partition of the arcs of G whose existence is guaranteed by Lemma 5.3. For each class, we choose a maximal antisymmetric subset (i.e., essentially an orientation). The graph H having the same node set as G but the union of all such antisymmetric sets as arcs is a regular graph (its degree is half of the degree of G). By symmetrization, the 1-factorization of H whose existence is guaranteed by Lemma 5.1 can be turned into a symmetric 2-factorization of G .

(2). One implication is straightforward (a graph of odd degree cannot be 2-factorable). For the other side, if G possesses a symmetric 1-factor, the graph obtained by deleting such a factor has even degree, and we can use part (1).

(3). The proof is similar to that of Lemma 5.2; the only relevant modification is in the construction of f and g —the requirements $f(x) \neq g(x)$, $f(f(x)) \neq x$ and $g(g(x)) \neq x$ must be fulfilled, so that the resulting 1-factor is antisymmetric (i.e., if a is in the factor \bar{a} is not). Then, in the construction of the transfinite sequence F_β one deletes from G not only the arcs in the F_α 's, but also their symmetric ones. This gives a symmetric 2-factorization of G of cardinality \mathfrak{c} . ■

The first part of the previous theorem is, of course, a generalized version of Petersen's theorem on 2-factorability of $2m$ -regular undirected graphs.

5.2 Regular Graphs

For every natural number d , the following theorem characterizes the covers of d -bouquets.

Theorem 5.3 The coverings of d -bouquets are exactly the d -regular graphs.

Proof. Each covering G of a d -bouquet is trivially d -regular by the local isomorphism property. On the other hand, Lemma 5.1 and Theorem 5.1 show that a d -regular graph covers a d -bouquet. ■

Clearly, by local isomorphism the left-to-right implication of the previous theorem is true for every cardinality \mathfrak{c} . Assuming well balancing, the reverse implication can be proved using Lemma 5.2:

Theorem 5.4 Let G be a well-balanced graph. Then G is \mathfrak{c} -regular iff it covers a \mathfrak{c} -bouquet.

Note that in the case of separated graphs we can just assume \mathfrak{c} -regularity. Using Theorem 5.2, we can easily extend the previous results to symmetric graphs without semi-edges:

Theorem 5.5 Let G be a symmetric graph without semi-edges. Then:

1. G covers symmetrically a $2m$ -bouquet iff it is $2m$ -regular;
2. G covers symmetrically a $(2m + 1)$ -bouquet iff it is $(2m + 1)$ -regular and possesses a symmetric 1-factor;
3. if G is well balanced, then G covers symmetrically a \mathfrak{c} -bouquet iff it is \mathfrak{c} -regular.

5.3 Schreier Graphs

Yet another characterization of graphs covering bouquets can be expressed in group-theoretical terms. Let Γ be a group, H a subgroup of Γ and $S \subseteq \Gamma$ a set of elements of Γ . The *Schreier graph* of Γ with respect to H and S is the graph having as nodes the right cosets of H , and an arc from Hg to Hh for each $s \in S$ such that $Hgs = Hh$. When S is closed by inversion, the resulting graph is naturally endowed with a symmetry.

Notice that it is natural to wonder whether a weaker definition, that is, that S be endowed with a self-inverse bijection $(\bar{}) : S \rightarrow S$ relating elements that are mapped to inverses by the right representation of Γ in the system of cosets of H (that is, $Hgs\bar{s} = Hg$ for all $g \in \Gamma$), would be more appropriate. However, our (restrictive) definition causes no loss of generality: in the case above, we simply consider the Schreier graph of $S_{[\Gamma:H]}$ with respect to the subgroup fixing a chosen node and the set of permutations S' induced by S ; the resulting graph is isomorphic to the original one, and moreover now the symmetry of S' is exactly inversion in $S_{[\Gamma:H]}$.

A Schreier graph is $|S|$ -regular, and it is connected iff $H\langle S \rangle = \Gamma$. Moreover, it is naturally coloured on the set S —the s -induced arc from Hg to Hgs is coloured by s ; such a graph is called the *Schreier colour graph* of Γ with respect to H and S . (When $H = 1$, we obtain the *Cayley (colour) graph* of Γ with respect to S .) Finally, we note that S can also be a *multiset* of elements of Γ , with obvious extensions, and we shall tacitly use this fact.

Theorem 5.6 The (symmetric) coverings of bouquets are exactly the (symmetric) Schreier graphs.

Proof. If G is the (symmetric) Schreier graph of Γ with respect to H and S , then we can build a covering onto a (symmetric) $|S|$ -bouquet by sending an arc that would be coloured by s in s . (The symmetry on the bouquet is the symmetry induced by inversion in Γ .)

Let now G be a (possibly symmetric) covering of a bouquet B , S_{N_G} the symmetric group on N_G and H the subgroup of S_{N_G} fixing a chosen node z . For each loop a of B consider the permutation $\pi_a \in S_{N_G}$ induced by oplifting a (i.e., $\pi_a(x) = t(x\tilde{a})$). Note that if a and b are exchanged by a symmetry of B , then $\pi_a = \pi_b^{-1}$.

We show that the Schreier graph of S_{N_G} with respect to H and $\{\pi_a \mid a \in A_B\}$ is isomorphic to G (note that the previous set could be really a multiset). The map on the nodes is obvious, as each coset of H in S_{N_G} is uniquely characterized by the element to which z is mapped, and the a -induced arc from $H\rho$ to $H\rho\pi_a$ is mapped to $\rho^{(z)}\tilde{a}$. By uniqueness of oplifting, this defines a graph isomorphism. ■

By combining the previous results, we also obtain that

Corollary 5.1 All d -regular graphs are Schreier graphs. All well-balanced \mathfrak{c} -regular graphs are Schreier graphs.

Corollary 5.2 Let G be a symmetric graph without semi-edges. Then, under any of the following hypotheses:

1. G is $2m$ -regular;
2. G is $(2m + 1)$ -regular and possesses a symmetric 1-factor;
3. G is \mathfrak{c} -regular and well balanced;

we have that G is a symmetric Schreier graph.

The results for the finite symmetric case are well known (see [14], where the authors remark that there are cubic graphs without a perfect matching, so not all symmetric $2m + 1$ -regular graphs are Schreier graphs).

In Table 1 and 2 we summarize the main results obtained in this section. Note that there are some gaps that are still to be filled—we do not know which regular graphs can cover a bouquet by a covering projection that is associated with an action; moreover, the classification for the symmetric case only applies to graphs without semi-edges. Note also that assuming well balancing, in the infinite degree case “Schreier” can be replaced with “regular”.

	Finite degree	Arbitrary degree
fibrations	irregular graphs	
fibrations assoc. with an action	node-transitive graphs	
fibrations assoc. with a free action	Cayley graphs	
coverings	regular graphs	Schreier graphs
coverings assoc. with an action	?	?
coverings assoc. with a free action	Cayley graphs	

Table 1: A classification of graphs fibred over bouquets.

	Finite degree	Arbitrary degree
symmetric coverings	$\left. \begin{array}{l} 2m\text{-reg. graphs} \\ 2m + 1\text{-reg. graphs with a symm. 1-factor} \end{array} \right\}$	Schreier graphs
symm. coverings assoc. with an action	?	?
symm. coverings assoc. with a free action	Cayley graphs	

Table 2: A classification of symmetric graphs without semi-edges fibred over bouquets.

The interesting lesson to be learned from the theorems above concerns the rôle of loops. In the classical treatment of finite undirected Schreier graphs we just mentioned there is an evident hiatus between graphs of even and odd degree, which shows up in Corollary 5.2. Essentially, all undirected graphs of even degree are Schreier, but this does not happen in the odd degree case—you need a 1-factor. The results of this section show that this hiatus is an artifact of the representation used, rather than a feature. If there are semi-edges, even a symmetric graph of even degree could need to possess a symmetric 1-factor to cover a bouquet (and thus be Schreier): as an example, consider the symmetric 2-regular graph with exactly two nodes and two semi-edges. Correspondingly, Theorem 5.1 and 5.6 do not exhibit special cases related to parity.

Another interesting consideration concerns the classical definition of covering between undirected graphs. There is no agreement in the literature about the definition of the lifting of a loop, and more generally about its very nature: should it be counted once or twice? In the first case the coverings of a bouquet are all regular graphs, in the second case all even degree regular graphs.

(Compare this fact with the very simple and general statement of Theorem 5.6, which just cannot be expressed in the language of undirected graphs.)

To answer this question, we must borrow some material from the next section. We think that the only reasonable mathematical answer to such a problem is to find a subcategory of \mathcal{S} , the category of symmetric graphs, which is equivalent to the category of undirected graphs with edge set represented as a multiset of unordered pairs of nodes and morphisms preserving adjacency. Indeed, such a category exists, and it is the full subcategory of \mathcal{S} induced by the graphs without loops that are not fixed by the symmetry (i.e., all loops are semi-edges); of course, in this equivalence every nonloop edge is mapped to pair of symmetric arcs, and every loop edge to an arc fixed by the symmetry (i.e., a semi-edge). Note that mapping loop edges to pairs of loops exchanged by the symmetry *would not work*, as, for instance, the automorphism group of the one-node, one-edge graph would contain just the identity in the undirected setting, and two different morphisms in the symmetric representation—the subcategory would not be full.

If we accept this viewpoint, we must inherit from \mathcal{S} the combinatorial definition of covering of undirected graph in the presence of loops: loop edges should be counted *once*, and should lift to a 1-factor (i.e., to a perfect matching). Note that in this way a perfect matching can also contain loops, and moreover every symmetric d -factor of the symmetric representation corresponds exactly to a d -factor in the classical sense.

This choice, however, would be in contrast with some literature, where undirected graphs have *two kinds of edges incident on a single vertex*, of degree two and one, respectively, the *loop edges* and the *semi-edges*. A semi-edge adjacent to x lifts to a perfect matching between the nodes in the fibre of x , while a loop edge lifts to a 2-factor. Semi-edges are usually introduced *a posteriori*, as an additional kind of edge, while loop edges are the standard “singleton edges” coming from the definition of an undirected graph (i.e., edges represented by the degenerate unordered pair given by the singleton $\{x\}$).

There is of course no mandatory choice: one can consider singleton edges as of degree one or two, and then add a special definition for an additional entity of degree two or one, respectively. However, if we keep in our mind the symmetric representation, there is a major problem in the way loop edges and semi-edges are handled above: loop edges of degree two should be “reversible”, in the intuitive sense that their two extremities should be permutable by an automorphism, in the same way one can permute two loop arcs exchanged by the symmetry. Clearly, the standard definition of undirected graph does not allow this, and this is the reason why we find more reasonable to consider undirected loop edges as counting once: when a clear distinction between the two kind of edges is required, symmetric graphs are the way to go.

6 A categorical standpoint

As we remarked in the introduction, the definition of graph fibration can be traced back to Grothendieck’s notion of fibration between categories (indeed, this seems to be the oldest ancestor of graph fibrations). Every graph has an associated category (built by the left adjoint to the forgetful functor sending a category to its base graph): objects are given by the nodes of the graph, while arrows are given by paths, with composition defined by concatenation; we usually denote with G both a graph and the free category it generates. Definition 2.1 can be simply restated as follows: $\varphi : G \rightarrow B$ is a fibration iff the induced functor $\varphi^* : G \rightarrow B$ is a (categorical) fibration [9], which turns out to be necessarily discrete. In the case of symmetric graphs, the natural free category is built by the left adjoint to the forgetful functor sending a category to its base graph endowed with the symmetry $\bar{f} = f^{-1}$: this time, arrows are given by paths quotiented with respect to the relation $a\bar{a} = \mathbf{1}$,

$\bar{a}a = \mathbf{1}$; again, we denote with G both a symmetric graph and the free category it generates (which, of course, turns out to be a groupoid), and Definition 2.3 can be restated as before.

Thus, one naturally expects graph fibrations to enjoy good categorical properties. Indeed, we shall see that many graph-theoretical constructions we have used in the previous sections can be naturally and elegantly described in categorical terms. Moreover, by proving that (op)fibrations are preserved by pullbacks we shall be able to give some results about common fibrations and coverings.

Graphs form a topos (i.e., a cartesian closed category with finite limits and a subobject classifier—see [20]) \mathcal{G} that can be handily described as the functor category $\mathbf{Sets}^{\mathcal{C}^{op}}$, where \mathcal{C} is the category with two objects N and A and two parallel arrows $s, t : N \rightarrow A$ between them. Analogously, symmetric graphs form a topos \mathcal{S} that can be described as $\mathbf{Sets}^{\mathcal{D}^{op}}$, where \mathcal{D} is built from \mathcal{C} by adding an involution on A satisfying obvious equations w.r.t. s and t . Both topoi are complete and cocomplete, as they are presheaf categories, and have been intensively studied [18, 33, 35, 34]. Note that the inclusion $\mathcal{C} \rightarrow \mathcal{D}$ induces a functor $\mathcal{S} \rightarrow \mathcal{G}$, which forgets the symmetry and has a left adjoint $\mathbf{Sym}(-)$ building the formal symmetrization whose elementary description was given in Section 2.

There is a very elegant characterization of categorical discrete (op)fibrations that can be easily carried over to the case of graphs (we thank Frank Piessens for bringing this fact to our attention): a morphism $\varphi : G \rightarrow B$ is a fibration iff the following square

$$\begin{array}{ccc} A_G & \xrightarrow{t} & N_G \\ \varphi \downarrow & & \downarrow \varphi \\ A_B & \xrightarrow{t} & N_B \end{array} \quad (1)$$

is a pullback (dually, φ is an opfibration iff the analogous square with t replaced by s is a pullback). Note that the square is simply half of the commutativity conditions of a graph morphism. This characterization makes it obvious that (symmetric) (op)fibrations are closed by composition, that is, there is a subcategory of \mathcal{G} that contains all graphs and (symmetric) (op)fibrations between them. Moreover, in the following commutative diagram

$$\begin{array}{ccccc} A_G & \xrightarrow{(\ulcorner)} & A_G & \xrightarrow{t} & N_G \\ \varphi \downarrow & & \varphi \downarrow & & \downarrow \varphi \\ A_B & \xrightarrow{(\urcorner)} & A_B & \xrightarrow{t} & N_B \end{array}$$

the right square is a pullback iff φ is a fibration, and the left square is a pullback iff φ commutes with the symmetries of G and B . If both things happen, then the whole square is a pullback, and this is true iff φ is also an opfibration (remember that $s = t \circ (\ulcorner)$). Thus, we have proved that

Proposition 6.1 A symmetric graph morphism is a fibration iff it is an opfibration. In particular, if it is an (op)fibration it is a covering projection as well.

It is interesting to remark that the construction of the universal fibration can be expressed by a very simple adjunction: consider the category \mathcal{G}_\bullet of *rooted graphs*, that is, graphs with a selected node and morphisms that preserve it, and the category $\mathbf{Sets}^{\omega^{op}}$ of (undirected) trees [17]. There is an obvious full and faithful functor $I : \mathbf{Sets}^{\omega^{op}} \rightarrow \mathcal{G}_\bullet$ that sends a tree to an in-tree of \mathcal{G}_\bullet having the root as selected node. The right adjoint to this functor builds for every rooted graph, that is, for every graph G and every node x , a tree \tilde{G}^x , and the counit $v_G^x : I(\tilde{G}^x) \rightarrow G$ satisfies the following

universal property: for every tree T and every graph morphism $\xi : I(T) \rightarrow G$, there is a unique morphism of trees $f : T \rightarrow \tilde{G}^x$ that makes the following diagram commute:

$$\begin{array}{ccc} & & I(\tilde{G}^x) \\ & \nearrow I(f) & \downarrow v_G^x \\ I(T) & \xrightarrow{\xi} & G \end{array}$$

In particular, by choosing T as a path we obtain a bijection between the nodes of $I(\tilde{G}^x)$ and the paths of G terminating at x , which gives back the construction used in Theorem 3.2. Note that, being I a full inclusion, one can essentially identify \mathbf{Sets}^{opp} with a subcategory of \mathcal{G}_\bullet , and just say that there is a morphism $v_G^x : \tilde{G}^x \rightarrow G$ such that every morphism from an in-tree to G (with selected node x) lifts uniquely through v_G^x .

Also Theorem 3.1 has a categorical nature: indeed, it just claims that v_G^x is the initial object of the comma category of fibrations of rooted graphs having base G (with selected node x). More explicitly, for each graph H with selected node y and each fibration $\varphi : H \rightarrow G$ such that $\varphi(y) = x$ there is a unique lifting of v_G^x along φ , as in the following diagram:

$$\begin{array}{ccc} & & H \\ & \nearrow & \downarrow \varphi \\ \tilde{G}^x & \xrightarrow{v_G^x} & G \end{array}$$

6.1 Pullbacks

The characterization given by diagram (1) allows us also to prove easily the graph counterpart of the classical theorem [9, Proposition 8.1.15] about pullbacks of fibrations:

Theorem 6.1 The pullback of a fibration along an arbitrary morphism is a fibration.

Proof. Consider the following pullback square in \mathcal{G}

$$\begin{array}{ccc} J & \longrightarrow & G \\ \psi \downarrow & & \downarrow \varphi \\ H & \xrightarrow{\eta} & B \end{array}$$

where $\eta : H \rightarrow B$ is an arbitrary graph morphism. By pulling back the square (1) along η , we obtain the following commuting cube in \mathbf{Sets}

$$\begin{array}{ccccc} & & A_G & \xrightarrow{t} & N_G \\ & \nearrow & \downarrow t & & \downarrow \varphi \\ A_J & \xrightarrow{t} & N_J & \nearrow & \\ \psi \downarrow & & \downarrow \varphi & & \downarrow \psi \\ & \nearrow \eta & A_B & \xrightarrow{t} & N_B \\ A_H & \xrightarrow{t} & N_H & \nearrow \eta & \end{array}$$

where the three vertical sides adjacent to φ are pullbacks; by the associativity theorem [8, Proposition 2.5.9], also the remaining side is a pullback, and thus ψ is a fibration. ■

A simple dual argument and the remark that in presheaf categories limits are computed pointwise leads to the following

Corollary 6.1 The pullback of an opfibration is an opfibration. The pullback of a (symmetric) covering projection is a (symmetric) covering projection.

Finally, the previous results allow us to relate common bases and common total graphs as follows:

Theorem 6.2 Let G and H be graphs fibred over the same graph B . Then there is a graph $J \subseteq G \times H$ fibred over G and H . In particular, if G and H are finite, loopless or separated, so is J . The same holds for (symmetric) coverings.

Proof. All morphisms in the following pullback

$$\begin{array}{ccc} J & \longrightarrow & G \\ \downarrow & & \downarrow \\ H & \longrightarrow & B \end{array}$$

where B is the common base of G and H , are fibrations by Theorem 6.1. To complete the proof, recall that J has an injection in $G \times H$, and that finite, loopless and separated graphs are closed by formation of products and subobjects. The analogous proof for (symmetric) coverings uses Corollary 6.1. ■

The fact that categorical fibrations are preserved by pullbacks cannot be used to prove Theorem 6.1: indeed, the existence of a categorical fibration into the free category generated by H does *not* imply the existence of a graph fibration inducing it. More generally, one has to be careful in translating properties of categorical fibrations to graphs: for instance, the projection $G \times H \rightarrow H$ is not generally a fibration. However, by noting that products preserve pullbacks one can easily show that

Theorem 6.3 If $\varphi : G \rightarrow B$, $\psi : H \rightarrow C$ are fibrations then $\varphi \times \psi : G \times H \rightarrow B \times C$ is a fibration. The same holds for (symmetric) coverings.

The previous theorems have several interesting consequences. Recall that Leighton's Theorem [19] states that two undirected finite graphs with the same universal covering have a common finite covering. We can immediately prove the analogous result for fibrations, as by Theorem 4.2 two graphs have the same universal total graphs iff they have the same minimum base, so Theorem 6.2 can be applied:

Corollary 6.2 Let G and H be graphs with the same universal total graphs. Then there is a graph $J \subseteq G \times H$ fibred over G and H . In particular, if G and H are finite, loopless or separated, so is J .

In force of our results about minimum covering bases, and with the same notation as in Section 5, we can also state that

Corollary 6.3 Let G and H regular graphs of finite degree having the same universal covering (i.e., having the same degree d). Then there is a graph $J \subseteq G \times H$ covering G and H . In particular, if G and H are finite, loopless or separated, so is J .

The proof now uses the coverings projections on a d -bouquet whose existence is guaranteed by Theorem 5.3. This result (which trivializes in the infinite case) cannot be obtained by extending Leighton's proof, for the latter strictly depends on the symmetry of the graphs involved. An example of such a pullback is given in Figure 7; note that nodes in the first column of the pullback are mapped to the central node of the five-node graph, while the nodes in the first row can be mapped to any of the four nodes of the other graph; the mapping of the remaining nodes is forced by the mapping of the arcs. By analogous techniques, we also obtain the following theorem, which was




Figure 7: A pullback of two covering projections.

essentially proved in the finite case by Sachs [29] working directly with factorizations:

Corollary 6.4 Let G and H be Schreier graphs of the same degree. Then there is a Schreier graph $J \subseteq G \times H$ covering G and H . If G and H are symmetric, finite, loopless or separated, so is J .

Proof. By Theorem 5.6, G and H cover (symmetrically) a bouquet, so we can apply Theorem 6.2. Note that the graph J covers (symmetrically) a bouquet, so it is a (symmetric) Schreier graph, by the other side of Theorem 5.6. ■

Theorem 6.3 yields also very simple proofs of some known facts: denoting with 1 the terminal object of \mathcal{G} (the loop), we have that $G \times K_2 \xrightarrow{1 \times i} G \times 1 \cong G$, so $G \times K_2$ (the *Kronecker double covering* of G) does cover G (symmetrically, if G is symmetric). This is immediate, as the unique morphism $i : K_2 \rightarrow 1$ is a symmetric covering, so its product with the identity gives rise to a symmetric covering as well. Analogously, denoting with U_n the unidirectional n -cycle we have the less well-known fact that $G \times U_n$ covers G , as the product $G \times U_n \xrightarrow{1 \times i} G \times 1 \cong G$ shows.

6.2 The category of fibrations over a given base

Let us denote with $\mathbf{Fib}(-, B)$ the category of fibrations with base B , and morphisms given by commuting triangles

$$\begin{array}{ccc} G & \xrightarrow{\eta} & H \\ \varphi \searrow & & \swarrow \psi \\ & B & \end{array}$$

with η an arbitrary morphism of \mathcal{G} . The following representation theorem mimics the analogous result for discrete (categorical) fibrations:

Theorem 6.4 $\mathbf{Fib}(-, B) \simeq \mathbf{Sets}^{B^{op}}$.

Proof. We prove the statement by constructing a functor $\Theta : \mathbf{Sets}^{B^{op}} \rightarrow \mathbf{Fib}(-, B)$ that will be shown to be an equivalence. To each functor $F : B^{op} \rightarrow \mathbf{Sets}$ we associate a graph $\Theta(F)$ as follows: the nodes of $\Theta(F)$ are the disjoint sum of the partial sections of F , while the arcs are given by the disjoint sum of (the graphs of) the functions between sections. More formally,

$$\begin{aligned} N_{\Theta(F)} &= \sum_{x \in N_B} F(x) = \{\langle x, e \mid x \in N_B, e \in F(x) \rangle\} \\ A_{\Theta(F)} &= \sum_{a \in A_B} F(t(a)) = \{\langle a, e \mid a \in A_B, e \in F(t(a)) \rangle\}. \end{aligned}$$

Adjacency is defined in the obvious way: $s(\langle a, e \rangle) = \langle s(a), F(a)(e) \rangle$ and $t(\langle a, e \rangle) = \langle t(a), e \rangle$. A map $\varphi : \Theta(F) \rightarrow B$ is then defined by $\varphi(\langle x, e \rangle) = x$ and $\varphi(\langle a, e \rangle) = a$. A straightforward calculation shows that φ is indeed a graph morphism, and more precisely a fibration.

Given a natural transformation $\xi : F \rightarrow G$, the map

$$\sum_{x \in N_B} \xi_x : N_{\Theta(F)} \rightarrow N_{\Theta(G)}$$

is the node component of a graph morphism $\Theta(\xi)$ whose arc component is given by

$$\langle a, e \rangle \xrightarrow{\Theta(\xi)} \langle a, \xi_{t(a)}(e) \rangle.$$

We note that

$$\begin{aligned} \Theta(\xi)(s(\langle a, e \rangle)) &= \Theta(\xi)(\langle s(a), F(a)(e) \rangle) = \langle s(a), \xi_{s(a)}(F(a)(e)) \rangle \\ &= \langle s(a), G(a)(\xi_{t(a)}(e)) \rangle = s(\langle a, \xi_{t(a)}(e) \rangle) = s(\Theta(\xi)(\langle a, e \rangle)), \end{aligned}$$

and similarly for the target map, so $\Theta(\xi)$ is a graph morphism; finally, the triangle

$$\begin{array}{ccc} \Theta(F) & \xrightarrow{\Theta(\xi)} & \Theta(G) \\ \varphi \searrow & & \swarrow \varphi' \\ & B & \end{array}$$

can be easily proved to be commutative, so Θ is a functor. Fidelity of Θ is immediate (recall that the node component of $\Theta(\xi)$ is the disjoint sum of the components of ξ), and fullness can be shown

by noting that a morphism $\eta : \Theta(F) \rightarrow \Theta(G)$ in $\mathbf{Fib}(-, B)$ must map nodes of the form $\langle x, e \rangle$ to nodes of the form $\langle x, e' \rangle$, so for each node x we can define $\xi_x(e) = e'$ when $\eta(\langle x, e \rangle) = \langle x, e' \rangle$; by a long but straightforward calculation suitably exploiting the lifting property, ξ turns out to be a natural transformation such that $\Theta(\xi) = \eta$.

To complete the proof, we just have to show that every object φ of $\mathbf{Fib}(-, B)$ is isomorphic to the image of an object F of $\mathbf{Sets}^{B^{op}}$; this can be easily done by defining $F(x) = \varphi^{-1}(x)$, and $F(a^{op})(y) = s(\tilde{a}^y)$ for all nodes y in the fibre of $s(a^{op}) = t(a)$. ■

Denoting with $\mathbf{Cov}(-, B)$ the obviously defined category of coverings of B , a proof absolutely analogous to that of Theorem 6.5 shows that

Theorem 6.5 $\mathbf{Cov}(-, B) \simeq \mathbf{Sets}^{\mathbf{Sym}(B)^{op}}$.

In the case of symmetric graphs, we can define categories of *symmetric* fibrations or coverings with base B , obtaining the following

Theorem 6.6 For a symmetric graph B , $\mathbf{SymFib}(-, B) \simeq \mathbf{SymCov}(-, B) \simeq \mathbf{Sets}^{B^{op}}$.

Note that the theorems above prove a fundamental fact—namely, that $\mathbf{Fib}(-, B)$ (as well as $\mathbf{Cov}(-, B)$ and $\mathbf{SymCov}(-, B)$) is a topos! (This does not happen in the categorical case as \mathbf{Cat} is not a topos.) Besides, they give a nice geometrical interpretation to presheaves over free categories.

The topos structure of $\mathbf{Fib}(-, B)$ will not be studied here, but we want to remark that representable functors over x of $\mathbf{Sets}^{B^{op}}$ are mapped by the equivalence above to the universal fibration $\nu_B^x : \tilde{B}^x \rightarrow B$, which gives still another categorical characterization of its construction. Analogously, in the case of coverings the representable functor over x of $\mathbf{Sets}^{\mathbf{Sym}(B)^{op}}$ is (mapped to) the universal *covering* at x , and if B is symmetric the representable functor over x of $\mathbf{Sets}^{B^{op}}$ is (mapped to) the universal *symmetric covering* at x .

Thus, the Yoneda Lemma, which gives a full embedding $B \hookrightarrow \mathbf{Sets}^{B^{op}}$ ($B \hookrightarrow \mathbf{Sets}^{\mathbf{Sym}(B)^{op}}$ in the case of coverings) where each node of B is mapped to its representable functor, shows in the case of (symmetric) coverings that the universal (symmetric) covering is the same along any connected component of B , since a walk in B is mapped to an isomorphism in $\mathbf{Sets}^{B^{op}}$, which is equivalent to an isomorphism in $\mathbf{SymCov}(-, B)$. In fact the Yoneda Lemma proves even more, namely that such an isomorphism commutes with the universal (symmetric) covering projections.

The object component of the equivalence described in Theorem 6.6 is well known: indeed, it essentially gives the representation of coverings (of undirected graphs) as graphs derived by *permutation voltage assignments* [14]. The categorical version, however, allows one not to choose a fixed set X for the fibres (and correspondingly the symmetric group on X for voltages), but rather to work in full generality; functoriality and naturality guarantee coherence.

The morphism component of the equivalence was not studied originally, but it is a cornerstone in Hofmeister's computation of the number of covering projections up to isomorphism [16]; he finds conditions on the permutation generating the covering that correspond exactly to the naturality squares. Thus, Theorems 6.4 and 6.5 are the natural generalization of the abovementioned results to fibrations and coverings of arbitrary graphs, and, correspondingly, one can use them to count isomorphism classes in $\mathbf{Fib}(-, B)$. In the next section we work out an example in this direction, using standard counting techniques.

As a final remark, we want to make the relation between graph bundles in the sense of Pisanski and Vrabec [27] and categorical fibrations explicit. Another fundamental topos of graphs is the topos \mathcal{R} of *symmetric reflexive graphs* [18]. In a symmetric reflexive graph every node x has an assigned *identity loop* ε_x that is preserved by morphisms and fixed by the symmetry (so it is a

semi-edge). More intuitively, we can identify such loops with nodes, and say that the morphisms in the topos are *degenerate*, in the sense that they can collapse an arc to a node (in fact, they just send the arc to the identity loop). Every symmetric reflexive graph has an associated category (in fact, a groupoid), built by the left adjoint to the forgetful functor sending a category to its base symmetric reflexive graph, in which the identity loops are exactly the identities of the category, and the symmetry is defined by $\overline{f} = f^{-1}$. Again, we define a fibration between symmetric reflexive graphs as a morphism that induces a fibration between the free categories they generate. Then, graph bundles in the sense of Pisanski and Vrabcic turn out to be fibrations of symmetric reflexive graphs, but of a particular kind, that we can characterize as follows: With each object F of the category $\mathcal{R}^{B^{op}}$, with B in \mathcal{R} , we can associate a fibration $G \rightarrow B$ in \mathcal{R} using a construction analogous to the proof of Theorem 6.4; more in detail, the total graph G is given by the disjoint union of the images $F(x)$ when x ranges through the nodes of B , enriched with an arc from $p \in F(x)$ to $q \in F(y)$ for every arc $a \in B$ from x to y such that $F(a)(p) = q$; the morphism from the total space to the base is now obvious, and such fibrations are exactly the bundles in the sense of Pisanski and Vrabcic. An axiomatic characterization of the categorical fibrations induced by such bundles is not currently known.

6.3 Counting minimal fibrations of the cycle

Let C_n be the bidirectional cycle on $n > 2$ nodes, having \mathbf{Z}_n as node set and an arc from x to $x \pm 1$. We want to count the number of nonisomorphic minimal fibrations of C_n , that is, the number of isomorphism classes of $\mathbf{Fib}(C_n, B)$, where B is the 2-bouquet (its arcs being denoted by a and b). By Theorem 6.4, every fibration in $\mathbf{Fib}(C_n, B)$ is equivalent to a functor $F : B^{op} \rightarrow \mathbf{Sets}$ that necessarily satisfies the equation $F(a^{op})(x) - x = x - F(b^{op})(x) \in \{-1, 1\}$ for every $x \in \mathbf{Z}_n$ (one has $|F(a^{op})(x) - x| = |F(b^{op})(x) - x| = 1$ by the definition of adjacency in C_n , and $F(a^{op})(x) \neq F(b^{op})(x)$, for C_n is separated). Such a functor is bijectively associated with the function (ambiguously denoted with) $F : \mathbf{Z}_n \rightarrow \{-1, 1\}$ satisfying $F(x) = F(a^{op})(x) - x$, so there are 2^n fibrations from C_n to B .

To apply Burnside's Lemma, we note that $\mathbf{Aut}(C_n)$ acts by precomposition on $\mathbf{Fib}(C_n, B)$ and that the orbits of this action are exactly the isomorphism classes we are to count. By the previous remarks, this is the same as counting the isomorphism classes of functions $\mathbf{Z}_n \rightarrow \{-1, 1\}$ under the (obviously induced) action of $\mathbf{Aut}(C_n)$. Note that the latter action is not the standard action of $\mathbf{Aut}(C_n)$ on a two-coloured cycle, so one cannot directly apply coloured counting techniques such as Pólya's Theorem or deBruijn's formula.

By definition of natural equivalence, an automorphism α of C_n fixes a fibration (represented as a functor $F : B^{op} \rightarrow \mathbf{Sets}$) exactly when the naturality equations $\alpha \circ F(a^{op}) = F(a^{op}) \circ \alpha$ and $\alpha \circ F(b^{op}) = F(b^{op}) \circ \alpha$ are satisfied; but the latter can be rewritten in the following simple form

$$\alpha(x \pm F(x)) = \alpha(x) \pm F(\alpha(x)) \quad (2)$$

using the function $F : \mathbf{Z}_n \rightarrow \{-1, 1\}$ we just defined. Note that $\mathbf{Aut}(C_n)$ has $2n$ elements, divided as follows:

- (i). n rotations (including the identity) $1, \rho, \rho^2, \dots, \rho^{n-1}$, defined by $\rho^k(x) = x + k$;
- (ii). symmetries with fixed points; if n is odd, then there is one symmetry σ_k for each node (the symmetry σ_k has only k as fixed point), defined by $\sigma_k(x) = 2k - x$; if n is even there are $n/2$ symmetries with two fixed points (say $\sigma_{0, n/2}, \sigma_{1, 1+n/2}, \dots, \sigma_{n/2-1, n-1}$), defined again by $\sigma_{k, k+n/2}(x) = 2k - x$;

- (iii). in the even case, there are also $n/2$ symmetries without fixed points $\tau_0, \tau_1, \dots, \tau_{n/2-1}$, defined by $\tau_k(x) = 2k - x + 1$.

Correspondingly, equation (2) provides the following conditions:

- (i). in order for ρ^k to fix F we must have $F(x) = F(x + k)$; thus exactly (n, k) values (by (n, k) we denote the *greatest common divisor* of n and k) of F can be independently chosen, amounting to $2^{(n,k)}$ fibrations fixed by ρ^k ;
- (ii). a symmetry σ with fixed points imposes the condition $F(x) = -F(2k - x)$, which is never satisfiable (just take $x = k$);
- (iii). for the remaining case, τ_k gives the constraint $F(x) = F(2k - x + 1)$; thus, exactly $n/2$ values of F can be independently chosen, amounting to $2^{\frac{n}{2}}$ fibrations.

Applying Burnside's Lemma we obtain that the number of orbits of the action of $\mathbf{Aut}(C_n)$, that is, isomorphism classes in $\mathbf{Fib}(C_n, B)$, is

$$\frac{1}{2n} \left(\sum_{k=0}^{n-1} 2^{(n,k)} + (n \text{ even})n2^{\frac{n}{2}-1} \right) = \frac{1}{2n} \sum_{k=0}^{n-1} 2^{(n,k)} + (n \text{ even})2^{\frac{n}{2}-2} \sim \frac{2^n}{2n}, \quad (3)$$

where we used Iverson's notation [12]: a (logical) formula enclosed in parenthesis takes value 1 when it is true, 0 otherwise.

For sake of completeness, we note that also $\mathbf{Aut}(C_n) \times \mathbf{Aut}(B)^{op}$ acts on $\mathbf{Fib}(C_n, B)$; in this case, the orbits are larger than the isomorphism classes of $\mathbf{Fib}(C_n, B)$, as two fibrations live in the same orbit if they differ by precomposition with an automorphism of C_n or by postcomposition with an automorphism of B , but we can still compute their number (more precisely, in this case the orbits are the isomorphism classes of the category having as objects fibrations and as morphisms commutative squares). For pairs of automorphisms with nontrivial second component ξ^{op} , equation (2) becomes

$$\alpha(x \pm F(x)) = \alpha(x) \mp F(\alpha(x)),$$

and we can easily extend our previous considerations as follows:

- (i). in order for $\langle \rho^k, \xi^{op} \rangle$ to fix F we must have $F(x) = -F(x + k)$; this equation is satisfiable only if $n/(n, k)$ is even, and in that case exactly (n, k) values of F can be independently chosen, amounting to $2^{(n,k)}$ fibrations;
- (ii). $\langle \sigma, \xi^{op} \rangle$ imposes the condition $F(x) = F(2k - x)$: this amounts to $2^{\frac{n}{2}+1}$ fibrations in the even case and to $2^{\frac{n+1}{2}}$ fibrations in the odd case;
- (iii). for the remaining case, $\langle \tau_k, \xi^{op} \rangle$ gives the constraint $F(x) = -F(2k - x + 1)$; thus, exactly $n/2$ values of F can be independently chosen, amounting to $2^{n/2}$ fibrations.

Summing up, in the even case we have

$$\frac{1}{4n} \left(\sum_{k=0}^{n-1} [1 + (n/(n, k) \text{ even})] 2^{(n,k)} + n2^{\frac{n}{2}+1} \right) = \frac{1}{4n} \sum_{k=0}^{n-1} [1 + (n/(n, k) \text{ even})] 2^{(n,k)} + 2^{\frac{n}{2}-1} \quad (4)$$

n	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
(3)	2	4	4	9	10	22	30	62	94	192	316	623	1096	2122	3856	7429	13798	26500
(4),(5)	2	4	4	8	9	18	23	44	63	122	190	362	612	1162	2056	3914	7155	13648

Table 3: Some values for the number of nonisomorphic minimal fibrations of C_n .

orbits, while in the odd case we obtain

$$\frac{1}{4n} \left(\sum_{k=0}^{n-1} 2^{(n,k)} + n2^{\frac{n+1}{2}} \right) = \frac{1}{4n} \sum_{k=0}^{n-1} 2^{(n,k)} + 2^{\frac{n-3}{2}}. \quad (5)$$

Both previous formulae are asymptotic to $2^n/4n$. This means that the number of orbits of the action of $\mathbf{Aut}(C_n)$ on minimal fibrations of an n -cycle is asymptotically twice the number of the orbits of the action of $\mathbf{Aut}(C_n) \times \mathbf{Aut}(B)^{op}$ (we give some exact values in Table 3), that is, ξ^{op} almost never fixes an orbit of $\mathbf{Aut}(C_n)$. Recalling that an automorphism β of B can be *lifted along a morphism* $\varphi : G \rightarrow B$ if there is an automorphism α of G such that $\alpha \circ \varphi = \varphi \circ \beta$, we can see that ξ can almost never be lifted along a minimal fibration of C_n . More precisely, ξ can be lifted along φ iff $\langle \alpha, \xi^{op} \rangle$ fixes φ for some automorphism α of C_n . Thus, our previous considerations show also that

Theorem 6.7 The nontrivial automorphism ξ of the 2-bouquet can almost never be lifted along a minimal fibration of C_n . More precisely,

$$\frac{|\{\varphi \in \mathbf{Fib}(C_n, B) \mid \xi \text{ can be lifted along } \varphi\}|}{|\mathbf{Fib}(C_n, B)|} = O(n/2^{\frac{n}{2}}).$$

We remark that the analogous problem for (symmetric) coverings is rather trivial: a slight modification of the previous techniques shows that

$$\begin{aligned} |\mathbf{Cov}(C_n, B)/\mathbf{Aut}(C_n)| &= |\mathbf{Cov}(C_n, B)/\mathbf{Aut}(C_n) \times \mathbf{Aut}(B)^{op}| = 1 + (n \text{ even}) \\ |\mathbf{SymCov}(C_n, B)/\mathbf{Aut}(C_n)| &= |\mathbf{SymCov}(C_n, B)/\mathbf{Aut}(C_n) \times \mathbf{Aut}(B)^{op}| = 1, \end{aligned}$$

where we assumed that the symmetry of B is nonidentical (in the other case, the result is 1 if n is even and 0 otherwise). The problem of computing the number of isomorphism classes of the category $\mathbf{Fib}(G, \widehat{G})$ for an arbitrary finite graph G is of course much more difficult, and will be pursued elsewhere.

7 Open problems

We conclude by formulating a series of open problems concerning fibrations of graphs; of course, the list given below has no claim of being exhaustive.

Problem 1 (Generalized bounds for Theorem 4.5) The theorem is true for (symmetric) coverings, but it is not known whether the bound given is tight in this case.

Problem 2 (Counting fibration-prime graphs) Given a natural number m , count the number of (strongly connected) graphs with m arcs that are fibration prime (analogously, one can fix the the number of nodes and a bound k for the number of parallel arcs). The same problem can be posed for (symmetric) coverings. Note that the knowledge of the asymptotic distribution of fibration (covering)-prime graphs would have immediate applications, as, for instance, it would allow one to estimate the probability of success of an anonymous election algorithm (see Section 2.5).

Problem 3 (Complete Table 1 and 2) Classify graphs that cover a bouquet *via* a projection associated with an action, and symmetric graphs with semi-edges covering symmetrically a bouquet.

Problem 4 (Classification of total graphs over B) Given a strongly connected fibration-prime graph B (or a set of such graphs), classify the graphs that are fibred over B . An analogous classification can be carried out for (symmetric) coverings, similarly to Section 5. An even more challenging question is the classification of graphs that cover *no prime covering other than B* .

Problem 5 (Counting fibrations) Count the number of nonisomorphic fibrations over a fixed graph B whose total space has a given number of nodes (or is isomorphic to another fixed graph G). In particular, count the isomorphism classes of $\mathbf{Fib}(G, \widehat{G})$ (for related results on undirected graph coverings, see [16]).

Problem 6 (Fibrations of reflexive graphs) An interesting theoretical problem is the study of fibrations in the topos of *reflexive* graphs, as in that case the fibrations induced between free categories are not necessarily discrete, so Definition 2.1 does not apply.

Problem 7 (Deciding nonemptiness of $\mathbf{Fib}(G, B)$) Given finite graphs G and B , it is decidable whether G is fibred over B . Study the complexity of this decision problem, in particular for fixed G or B .

References

- [1] *Revêtements étales et groupe fondamental. Séminaire de Géométrie Algébrique du Bois Marie 1960–1961 (SGA 1). Dirigé par Alexandre Grothendieck. Augmenté de deux exposés de M. Raynaud.*, volume 224 of *Lecture Notes in Mathematics*. Springer–Verlag, 1971.
- [2] Dana Angluin. Global and local properties in networks of processors. In *Proc. 12th Symposium on the Theory of Computing*, pages 82–93, 1980.
- [3] Dwight R. Bean. Effective coloration. *J. Symbolic Logic*, 41(2):469–480, 1976.
- [4] Claude Berge. *Graphs and Hypergraphs*. North–Holland, Amsterdam, 1973.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Francis Borceux. *Handbook of Categorical Algebra 1*, volume 50 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, 1994.
- [9] Francis Borceux. *Handbook of Categorical Algebra 2*, volume 51 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, 1994.
- [10] Younki Chae, Jin Ho Kwak, and Jaeun Lee. Characteristic polynomials of some graph bundles. *J. Korean Math. Soc.*, 30(1):229–249, 1993.

- [11] Dragoš M. Cvetković, Michael Doob, and Horst Sachs. *Spectra of Graphs*. Academic Press, 1978.
- [12] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Addison–Wesley, second edition, 1994.
- [13] John W. Gray. Fibred and cofibred categories. In *Proceedings of the conference on categorical algebra: La Jolla 1965*, pages 21–83. Springer–Verlag, 1966.
- [14] Jonathan L. Gross and Thomas W. Tucker. *Topological Graph Theory*. Series in Discrete Mathematics and Optimization. Wiley-Interscience, 1987.
- [15] Alexandre Grothendieck. Technique de descente et théorèmes d’existence en géométrie algébrique, I. Généralités. Descente par morphismes fidèlement plats. *Seminaire Bourbaki*, 190, 1959–1960.
- [16] Michael Hofmeister. Isomorphisms and automorphisms of graph coverings. *Discrete Math.*, 98:175–183, 1991.
- [17] Stefano Kasangian and Sebastiano Vigna. The topos of labelled trees: A categorical semantics for SCCS. *Fund. Inform.*, 32:27–45, 1997.
- [18] F. William Lawvere. Qualitative distinctions between some toposes of generalized graphs. In *Categories in computer science and logic (Boulder, CO, 1987)*, volume 92 of *Contemp. Math.*, pages 261–299. American Mathematical Society, 1989.
- [19] Frank Thomson Leighton. Finite common coverings of graphs. *J. Combin. Theory Ser. B*, 33:231–238, 1982.
- [20] Colin McLarty. *Elementary Categories, Elementary Toposes*. Number 21 in Oxford Logic Guides. Oxford University Press, 1992.
- [21] Robin Milner. *Communication and Concurrency*. International Series in Computer Science. Prentice Hall, 1989.
- [22] Abbe Mowshowitz. The adjacency matrix and the group of a graph. In *New directions in the theory of graphs (Proc. Third Ann Arbor Conf. on Graph Theory, Univ. of Michigan, Ann Arbor, Mich., 1971)*, pages 129–148. Academic Press, 1973.
- [23] Nancy Norris. Universal covers of graphs: Isomorphism to depth $n - 1$ implies isomorphism to all depths. *Discrete Appl. Math.*, 56:61–74, 1995.
- [24] Oynstein Ore. *Theory of Graphs*, volume XXXVIII of *American Mathematical Society Colloquium Publications*. American Mathematical Society, 1962.
- [25] M. Petersdorf and Horst Sachs. Über Spektrum, Automorphismengruppe und Teiler eines Graphen. *Wiss. Z. Techn. Hochsch. Ilmenau*, 15:123–128, 1969.
- [26] M. Petersdorf and Horst Sachs. Spektrum und Automorphismengruppe eines Graphen. In *Combinatorial Theory and its Applications, III (Proc. Colloq., Balatonfüred, 1969)*, pages 891–907. North–Holland, 1970.
- [27] Tomaž Pisanski, John Shawe–Taylor, and Jože Vrabec. Edge-colorability of graph bundles. *J. Combin. Theory Ser. B*, 35(1):12–19, 1983.
- [28] Gert Sabidussi. On a class of fixed-point free graphs. *Proc. Amer. Math. Soc.* 9, 5:800–804, 1958.
- [29] Horst Sachs. Simultane Überlagerung gegebener Graphen. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 9:415–427, 1965.
- [30] Horst Sachs. Über Teiler, Faktoren und charakteristische Polynome von Graphen. I. *Wiss. Z. Techn. Hochsch. Ilmenau*, 12:7–12, 1966.
- [31] Allen J. Schwenk. Review #8897. *Math. Rev.*, 46:1538, 1973.
- [32] Allen J. Schwenk. Computing the characteristic polynomial of a graph. In Ruth A. Bari and Frank Harary, editors, *Graphs and Combinatorics*, volume 406 of *Lecture Notes in Mathematics*, pages 153–172. Springer–Verlag, 1974.

- [33] John Shrimpton. Some groups related to the symmetry of a directed graph. *J. Pure Appl. Algebra*, 72(3), 1991.
- [34] Sebastiano Vigna. A guided tour in the topos of graphs. Technical Report 199-97, Università di Milano, Dipartimento di Scienze dell'Informazione, 1997.
- [35] Oswald Wyler. *Lecture Notes on Topoi and Quasitopoi*. World Scientific, 1991.
- [36] Masafumi Yamashita and Tiko Kameda. Computing functions on asynchronous anonymous networks. *Math. Systems Theory*, 29(4):331–356, 1996. See also [38].
- [37] 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.
- [38] Masafumi Yamashita and Tiko Kameda. Erratum to computing functions on asynchronous anonymous networks. *Theory of Computing Systems (formerly Math. Systems Theory)*, 31(1), 1998.