

On the Complexity of Deciding Sense of Direction*

Paolo Boldi[†]

Sebastiano Vigna[†]

Abstract

In this paper we prove that deciding whether a distributed system (represented as a coloured digraph with n nodes) has weak sense of direction is in AC^1 (using n^6 processors). Moreover, we show that deciding sense of direction is in P. Our algorithms can also be used in order to decide in AC^1 whether a coloured graph is a Cayley colour graph.

1 Introduction

The theory of distributed computing aims at understanding the very nature of cooperation in a distributed environment, where processing is carried on by agents with autonomous computing abilities, which can communicate by exchanging messages along communication links. The topological structure of such systems can be described by a graph, with nodes representing agents and arcs representing links. Each node has a local (partial) view of the system, and it associates a different label (colour) to each of its incident links; in general, the colour assigned to a certain link by a node may differ from the one assigned to the same link by the partner, because colours are assigned locally.

The solution to many problems in a distributed system can be greatly simplified by using colourings with special properties. In this paper we analyse from the point of view of complexity theory a property known as *(weak) sense of direction*. Suppose that a node wants to know whether two messages have been sent by the same source, and assume that each message is attached a string listing the colours of the links it has been received from, along the route it has covered. Since each node assigns a different colour to each of its incident arcs, if every agent has a complete view of the system (and of every local colouring as well), the problem can be solved by an easy local algorithm which uses a linear backtracking technique to single out the source. But is there a way to solve the problem *globally*, with a unique function which also abstracts partially from the overall knowledge of the system? The possibility for this to happen clearly depends on some property of the colouring.

This property of global consistency has been known and used informally for a long time (see, for instance, [LMW86, AvLSZ89, San84]), but was not formalized until [FMS95]; this

*A preliminary version of this paper appeared in the Proceedings of the 2nd Colloquium on Structural Information and Communication Complexity, Lefteris M. Kirousis and Evangelos Kranakis, eds., Carleton University Press, 1996.

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

last definition, now widely accepted, allowed a systematic study of this notion. In the present paper, we shall characterize for the first time weak sense of direction in a combinatorial manner, and show that it can be decided efficiently in parallel; namely, we shall prove that the presence of weak sense of direction can be decided in logarithmic time using n^6 processors.

However, having weak sense of direction is in general not enough for practical purposes: what one usually wants is the possibility of coding the identity of the source locally in each node, without using a complete path to decode it. In other words, each node should be able to infer the local name of the source simply using the colour of the link from which the message has arrived, and the name under which the next-to-last node on the path knows the source. This property is known as *sense of direction*: again, we shall characterize it combinatorially, and we will show that it is decidable in sequential polynomial time.

The upper bounds we derive are the first nontrivial ones for these decision problems, even though the algorithms we propose are not practical (due to the high number of processors, in the first case, and to the high degree of the polynomial bound, in the second case). Note however that an $O(n^4)$ sequential algorithm for weak sense of direction can be easily derived from the results of Section 5 by representing the graph as an adjacency list, and using Tarjan's algorithm for strongly connected components [Tar72].

In the last part of the paper, by exploiting the tight connection between weak sense of direction and Cayley graphs, we shall prove that the problem of recognizing Cayley colour graphs is in AC^1 . Finally, we shall discuss some open problems related to this area.

2 Definitions

A *directed graph* (or, in short, a graph) G is given by a set V of n vertices and a set $A \subseteq V \times V$ of arcs. We write $P[x, y] \subseteq A^*$ for the set of paths from the vertex x to the vertex y .

An (*arc*) *colouring* of a graph G is a function $\lambda : A \rightarrow \mathcal{L}$, where \mathcal{L} is a finite set of colours. We say that λ is *deterministic* iff

$$\lambda(\langle x, y \rangle) = \lambda(\langle x, z \rangle) \implies y = z$$

i.e., if the automaton described by the transition graph G with colouring λ is deterministic. We shall frequently shift between graph-theoretic and automata-theoretic concepts when talking about a coloured graph.

Our (coloured) graphs will be always represented by (coloured) adjacency matrices, i.e., matrices $n \times n$ such that the entry indexed by $x, y \in V$ will contain 0 if no arc connects x to y ; otherwise, it will contain a positive integer representing the colour of the arc (or 1, if the graph has no colouring).

Given a graph G deterministically coloured by λ (we shall omit in the future to mention that λ is deterministic; all colourings in this paper are such), let

$$L(x, y) = \{\lambda^*(\pi) \mid \pi \in P[x, y]\},$$

where $\lambda^* : A^* \rightarrow \mathcal{L}^*$ is the pointwise lifting of λ . In other words, $L(x, y)$ is the *language*

recognized by G when x is the initial state and y is the final state. For all $I \subseteq V^2$ let

$$L_I = \bigcup_{(x,y) \in I} L(x, y)$$

(of course, $L(x, y) = L_{\{(x,y)\}}$). Notice that $\varepsilon \in L(x, x) \neq \emptyset$.

A *local naming* for G is a family of injective functions $\beta = \{\beta_x : V \rightarrow \mathcal{N}\}_{x \in V}$, with \mathcal{N} a finite set, called the *name space*. Intuitively, each node x of G gives to each other node y a name $\beta_x(y)$ taken from the name space. Since we require injectivity, we have that necessarily $|\mathcal{N}| \geq n$. We shall also write $\beta : V \times V \rightarrow \mathcal{N}$ for the “unindexed” local naming.

Given a coloured graph endowed with a local naming, a function $f : L_{V^2} \rightarrow \mathcal{N}$ is a *coding function* iff

$$\forall x, y \in V \quad \forall \pi \in P[x, y] \quad f(\lambda^*(\pi)) = \beta_x(y).$$

A coding function translates the colouring of the path along which two nodes x, y are connected into the name which x gives to y . Note that while the resulting name is *local* (i.e., x and z might choose different element of the name space for the same node y), the coding function is *global* (i.e., it is the same for all nodes).

A colouring λ is a *weak sense of direction* for a graph G iff for some local naming there is a coding function¹. We shall also say that a coloured graph *has* weak sense of direction, or that λ *gives* weak sense of direction to G .

As an example, consider a $p \times q$ torus where the links have the standard “compass” colouring (North/South/East/West). Each node of coordinates (i, j) (where $i \in \mathbf{Z}_p, j \in \mathbf{Z}_q$) gives to a node of coordinates (h, k) the local name $(i - h, j - k)$. If a message arrives through a path coloured by v , the receiver knows the sender under the local name $f(v) = (\#_N(v) - \#_S(v), \#_E(v) - \#_W(v))$, where $\#_N(v)$ is the number of occurrences of the “North” colour in v , and so on.

Finally, a *decoding function* is a map $h : \mathcal{L} \times \mathcal{N} \rightarrow \mathcal{N}$ which satisfies

$$\forall (x, y) \in A \quad \forall z \in V \quad \forall \pi \in P[y, z] \quad h(\lambda((x, y)), f(\lambda^*(\pi))) = \beta_x(z).$$

A decoding function translates the name given by y to z into the name given by x to z knowing only the colour of the arc connecting x and y .

A colouring λ is a *sense of direction* for a graph G iff for some local naming there is a coding function and a decoding function. Our previous example has also trivially sense of direction (just add the contribution of the last colour).

The model of computation we use is a *Common CRCW PRAM* (i.e., concurrent reads and writes are allowed, and all the processors participating to a concurrent write must write the same value). We shall denote with AC^k the class of problems which are solvable in time $O((\log n)^k)$ using a polynomial number of processors (the classes AC^k were originally defined using boolean circuits, but they can be equivalently characterized in terms of parallel complexity; see [KR90]).

¹Elsewhere weak sense of direction has been defined in a slightly different way, by considering only non-empty paths. It is easy to check that the algorithms described here can be immediately adapted to that definition, just by assuming $\varepsilon \notin L(x, x)$ and, consequently, performing a transitive (non reflexive) closure of the matrix M in Proposition 1.

We remark that throughout this paper $|\mathcal{L}|$ is polynomially bounded in n . In particular, we can restrict without loss of generality to situations in which $|\mathcal{L}| \leq n^2$, for no more than n^2 colours can be actually “used” by λ .

Moreover, the notions of equivalence relation, boolean matrix and partition will be used interchangeably. Thus, we shall indifferently write $x R y$, $R(x, y) = 1$ or $x, y \in I \in R$. If R and S are equivalence relations, we shall write $R \leq S$ whenever R is finer than or equal to S as an equivalence relation, i.e., if $R \subseteq S$ (we shall also say that S is coarser than R). We shall use the same notation, with the same meaning, for the associated partitions and boolean matrices. Composition of relations (or, equivalently, product of matrices) will be denoted by juxtaposition, and reflexive/transitive closure of R by R^* .

3 Some remarks on notation

Our notation is rather different from the one which can be found in [FMS95]. In this section we show the rationale behind our choices, and we state the relations with the original notation. The reader unacquainted with [FMS95] may want to skip this section.

Instead of using undirected, connected graphs with a colouring which is node-dependent, we consider general coloured directed graphs², in which an arc from x to y represents the existence of a link *from* y *to* x ; this apparently unnatural convention makes our presentation much simpler³, and it allows us to use a more standard way of colouring graphs, which highlights naturally the connections with automata and regular languages. Our condition of *determinism* corresponds exactly to *local orientation* in [FMS95]. Moreover, our description of a graph makes immediately evident that the path-colouring function Λ_x used in [FMS95] is not really depending on x , so we denote it in a more standard fashion using the notation λ^* .

Note also that our definition of (de)coding function corresponds to the definition of a *consistent (de)coding function* of [FMS95]. We dropped the adjective “consistent” when referring to (de)coding functions, since the “inconsistent” version is never used in this paper. The original definition was stated in terms of partial functions; as far as f is concerned, we found simpler to specify explicitly its domain (L_{V^2}); h has been instead turned into a total function which has some “irrelevant” values (as opposed to “undefined”).

4 Coding functions and internal monodromy

We shall now give a purely combinatorial condition on the colouring of a graph which will be proved equivalent to being a weak sense of direction. From now onwards we shall work with finite graphs, and stick to the notation of Section 2 without further remarks.

²It is easy to extend our results to graphs with parallel arcs if the set of arcs between two vertices is represented by a vector of integers, one for each colour.

³The case studied in [FMS95] can be recovered by considering graphs which are symmetric, loopless, and strongly connected.

Definition 1 Given a set X , a partition Π of X^2 is said to be (*internally*) *monodrome* iff its elements are graphs of partial endofunctions of X , i.e., iff

$$\forall I \in \Pi \quad \langle x, y \rangle, \langle x, z \rangle \in I \implies y = z.$$

The previous definition will be intensively used throughout the paper. Note that any equivalence relation which is finer than a monodrome one is also monodrome.

Theorem 1 Let T be the equivalence relation on V^2 defined as the transitive closure of

$$\langle x, y \rangle \sim \langle x', y' \rangle \iff L(x, y) \cap L(x', y') \neq \emptyset.$$

Then λ is a weak sense of direction iff T is internally monodrome.

The proof of the theorem is preceded by two lemmata. Intuitively, each element I of T will be an element of the name space. The name under which x knows y is precisely the (unique) partial function $I \in T$ such that $\langle x, y \rangle \in I$. The idea behind the proof is that whenever the same string v of colours appears on two different paths of the graph, say from x to y and from x' to y' , then necessarily $\beta_x(y) = f(v) = \beta_{x'}(y')$. This set of constraints, which must not violate injectivity of the local namings, is resumed by the relation T .

Lemma 1 Let G be a graph with weak sense of direction given by a colouring λ , and let β and f be the corresponding local naming and coding function; then, there exists a monodrome partition $\Pi \geq T$ such that $|\Pi| \leq |\mathcal{N}|$ (recall that \mathcal{N} is the codomain of β).

Proof. We will show that the non-empty fibres of β (i.e., the non-empty counterimages of singletons) form a monodrome partition Π of V^2 which is coarser than T (and, of course, $|\mathcal{N}| \geq |\text{im}(\beta)| = |\Pi|$). Monodromy is trivial, as it is equivalent to injectivity of the naming functions β_x . Now suppose $\langle x, y \rangle, \langle x', y' \rangle \in I \in T$; then there is a chain

$$\langle x, y \rangle = \langle x_1, y_1 \rangle, \langle x_2, y_2 \rangle, \dots, \langle x_n, y_n \rangle = \langle x', y' \rangle$$

such that $w_i \in L(x_i, y_i) \cap L(x_{i+1}, y_{i+1}) \neq \emptyset$ for $1 \leq i \leq n-1$. But this means that $\beta_{x_i}(y_i) = f(w_i) = \beta_{x_{i+1}}(y_{i+1})$, which implies that $\beta_x(y) = \beta_{x'}(y')$, i.e., $\langle x, y \rangle, \langle x', y' \rangle \in J \in \Pi$. ■

Lemma 2 For every monodrome partition $\Pi \geq T$ there exists a local naming β with name space Π and coding function $f : \mathcal{L}^* \rightarrow \Pi$.

Proof. First of all, notice that the L_I 's form a partition of L_{V^2} when I ranges over T . Indeed, if $w \in L_I \cap L_J$ then $w \in L(x, y) \cap L(x', y')$ for some $\langle x, y \rangle \in I, \langle x', y' \rangle \in J$. But then necessarily $I = J$.

Let now $\beta_x(y)$ be the unique element of Π containing $\langle x, y \rangle$, and let $f(w)$ be the unique $I \in \Pi$ such that $w \in L_I$. Then, β is a local naming and f is a coding function; indeed:

- If $\beta_x(y) = \beta_x(z)$ then there is an element of Π including $\langle x, y \rangle$ and $\langle x, z \rangle$. But the definition of monodrome partition implies $y = z$.

- If $\pi \in P[x, y]$ and $\lambda^*(\pi) = w$ then $w \in L(x, y)$. Now, if $I \in \Pi$ is the unique element of Π which contains $\langle x, y \rangle$, then

$$f(\lambda^*(\pi)) = f(w) = I = \beta_x(y). \blacksquare$$

The proof of Theorem 1 is now trivial. If G has weak sense of direction, T is the refinement of a monodrome partition, and it is thus monodrome. On the other hand, by taking $\Pi = T$ in Lemma 2 we obtain a weak sense of direction from a graph G with T monodrome.

Note that even though monodromy and weak sense of direction are equivalent, it could still be the case that monodrome partitions coarser than T do not give the best possible (i.e., smallest) name spaces. Nonetheless, Lemma 1 guarantees that optimal naming is equivalent to optimal partitioning.

5 The decision problem for weak sense of direction

Using the results of Section 4 we shall now establish that deciding whether a given coloured graph G with n nodes has weak sense of direction is in AC^1 (using n^6 processors).

5.1 Checking monodromy

In this section we will show how to build T using reflexive/transitive closures of suitable matrices.

Proposition 1 Let G be a graph with colouring λ . Let M be the $n^2 \times n^2$ boolean matrix defined by

$$M(\langle x, x' \rangle, \langle y, y' \rangle) = 1 \iff \langle x, y \rangle, \langle x', y' \rangle \in A \wedge \lambda(\langle x, y \rangle) = \lambda(\langle x', y' \rangle).$$

Then,

$$L(x, y) \cap L(x', y') \neq \emptyset \iff M^*(\langle x, x' \rangle, \langle y, y' \rangle) = 1.$$

Intuitively, the matrix M is the non-coloured adjacency matrix of the *product graph*⁴ $G \times G$, which has an arc coloured by α between the nodes $\langle x, x' \rangle$ and $\langle y, y' \rangle$ iff G has arcs coloured by α between x, y and between x', y' .

Proof. $L(x, y) \cap L(x', y') \neq \emptyset$ happens iff there is a string $w = \alpha_1 \alpha_2 \cdots \alpha_n$ such that $x \cdot w = y$ and $x' \cdot w = y'$,⁵ i.e., iff there are paths

$$\langle x, x \cdot \alpha_1 \rangle \langle x \cdot \alpha_1, x \cdot \alpha_1 \alpha_2 \rangle \cdots \langle x \cdot \alpha_1 \alpha_2 \cdots \alpha_{n-1}, x \cdot w = y \rangle$$

and

$$\langle x', x' \cdot \alpha_1 \rangle \langle x' \cdot \alpha_1, x' \cdot \alpha_1 \alpha_2 \rangle \cdots \langle x' \cdot \alpha_1 \alpha_2 \cdots \alpha_{n-1}, x' \cdot w = y' \rangle.$$

⁴In the sense of topos theory, i.e., in the presheaf category of coloured graphs.

⁵We are again considering G as an automaton, and we are denoting by the dot operator the right action of the alphabet \mathcal{L} on states; in other words, $x \cdot \alpha$ is the state reached from x along the unique arc coloured by α , if it exists.

By the definition of M , $M(\langle z, z' \rangle, \langle z \cdot \alpha, z' \cdot \alpha \rangle) = 1$ whenever both $z \cdot \alpha$ and $z' \cdot \alpha$ are defined. Thus, $L(x, y) \cap L(x', y') \neq \emptyset$ is equivalent to reachability (i.e., reflexive/transitive closure) of $\langle y, y' \rangle$ from $\langle x, x' \rangle$ in the graph having M as adjacency matrix. ■

The previous proposition implies that $T = N^*$, where N is the matrix obtained from M^* using the following permutation:

$$N(\langle x, y \rangle, \langle x', y' \rangle) = M^*(\langle x, x' \rangle, \langle y, y' \rangle).$$

Once the matrix T has been computed, the monodromy condition is written as follows:

$$\forall x, y, z \in V \quad T(\langle x, y \rangle, \langle x, z \rangle) = 1 \implies y = z$$

As a matter of fact, this condition has a simple formulation in terms of submatrices: if we divide T into n^2 submatrices of size $n \times n$, the n submatrices on the diagonal have to be identities.

5.2 Weak sense of direction is in AC¹

We are now going to show that the operations described in the previous section can be realized using a logarithmic number of steps, and n^6 processors. We firstly give a formal definition of our problem:

Problem 1 WEAK SENSE OF DIRECTION

Instance: A graph G with colouring λ .

Question: Is λ a weak sense of direction?

Theorem 2 WEAK SENSE OF DIRECTION is in AC¹.

Proof. Given a coloured graph G with n nodes, we can compute the matrix M (using the notation of the previous section) in constant time using n^4 processors (i.e., one processor per entry of M). Then we can compute M^* in logarithmic time using n^6 processors (see [KR90]). But N can be computed from M using n^4 processors and constant time, after which $T = N^*$ can be computed again in logarithmic time using n^6 processors. Monodromy of T can be trivially checked in constant time using n^3 processors. ■

We state also the following result:


Theorem 3 The local naming β associated to T can be computed in constant time using n^6 processors.

Proof. Using n^6 processors and a constant number of steps we can compute for each row of N^* the index of the column of the first element which is nonzero. This number identifies uniquely each equivalence class (of course, $\beta_x(y)$ is exactly the number associated to the row $\langle x, y \rangle$). ■

6 Decoding functions and left regularity

We now attack the problem of characterizing coloured graphs which have sense of direction, i.e., which have a decoding function besides a local naming and a coding function.

Note that weakness is not a “fake” notion: there are coloured graphs which have weak sense of direction, but do not have a decoding function. An example follows:



Because of the structure of T , we are forced to set $f(bc) = f(d)$ for any coding function f . But as soon as we suppose the existence of a decoding function h , we obtain

$$\beta_4(5) = f(g) = \beta_0(1) = h(a, f(bc)) = h(a, f(d)) = \beta_2(3) = f(h) = \beta_4(6),$$

which is clearly impossible because β_4 is injective. Note that this example can be easily modified in order to obtain a symmetric connected graph with the same property (just add all arcs which are necessary assigning to each new arc a new colour).

Definition 2 Let G be a graph with colouring λ . A partition Π of V^2 is said to be *left regular* iff

$$\forall I \in \Pi \quad \forall \alpha \in \mathcal{L} \quad \exists J \in \Pi \quad \alpha L_I \cap L_{V^2} \subseteq L_J.$$

If a coloured graph admits a left regular monodrome partition $\Pi \geq T$, the language L_I ($I \in \Pi$), when prefixed with α , gives a sublanguage αL_I of some L_J , $J \in \Pi$. In other words, if two strings v, w live in the same language L_I , they cannot be “separated” by an α -prefixing. Note that if $\alpha L_I \cap L_{V^2} \neq \emptyset$, then J is unique.

Theorem 4 A graph G with colouring λ has sense of direction iff it admits a left regular monodrome partition $\Pi \geq T$.

Proof. If G admits a left regular monodrome partition $\Pi \geq T$, we set up f and β as in Lemma 2, and we define $h(\alpha, f(v))$ for $\alpha \in \mathcal{L}$, $v \in L_I$ as the unique $J \in \Pi$ such that $\alpha L_I \cap L_{V^2} \subseteq L_J$ (the choice of v is, of course, irrelevant). Trivially, this definition makes h into a decoding function.

Consider now a graph G for which the colouring λ is a sense of direction. Recalling the notation of Lemma 1, we write Π for the monodrome partition obtained by considering the non-empty fibres of β . Thus, we can identify elements in Π with elements of the image of β . We just have to prove that Π is left regular.

Given $\alpha \in \mathcal{L}$ and $I \in \Pi$, we consider strings $w_1, w_2 \in L_I$ such that $\alpha w_1, \alpha w_2 \in L_{V^2}$. Thus, we have vertices x_1, x_2 such that $x_1 \cdot \alpha w_1$ and $x_2 \cdot \alpha w_2$ are both defined. Then, we have

$$\beta_{x_1}(x_1 \cdot \alpha w_1) = h(\alpha, f(w_1)) = h(\alpha, f(w_2)) = \beta_{x_2}(x_2 \cdot \alpha w_2),$$

so each string in $\alpha L_I \cap L_{V^2}$ belongs to the same L_J , $J \in \Pi$. ■

The problem we face when using left regularity in order to check for the existence of a decoding function is that in principle we should check *all* possible monodrome partitions coarser than T . This clearly would lead to a combinatorial explosion. The next few theorems will be helpful in reducing the computational burden of Definition 2.

Definition 3 For each $\alpha \in \mathcal{L}$ let the relation $C_\alpha \subseteq V^2 \times V^2$ be defined by

$$\langle x, y \rangle C_\alpha \langle x', y' \rangle \iff \alpha L(x, y) \cap L(x', y') \neq \emptyset.$$

For each relation $R \subseteq V^2 \times V^2$ and for each $w = \alpha_1 \alpha_2 \cdots \alpha_n \in \mathcal{L}^*$ we set also

$$(CR)_w = C_{\alpha_1} R C_{\alpha_2} R \cdots C_{\alpha_n} R.$$

The operator $(-)^{C_\mathcal{L}}$, which transforms relations on V^2 , is defined by

$$R^{C_\mathcal{L}} = \left(\sum_{w \in \mathcal{L}^*} ((CR)_w)^T R (CR)_w \right)^*.$$

We observe that $(-)^{C_\mathcal{L}}$, when applied to a *symmetric* relation, always gives an *equivalence* relation. Moreover, it is monotone (i.e., $R \leq S$ implies $R^{C_\mathcal{L}} \leq S^{C_\mathcal{L}}$), and $R \leq R^{C_\mathcal{L}}$.

Proposition 2 A partition Π of V^2 coarser than T is left regular iff $\Pi^{C_\mathcal{L}} \leq \Pi$.

Proof. We have that

$$\begin{aligned} \Pi^{C_\mathcal{L}} \leq \Pi &\iff \left(\sum_{w \in \mathcal{L}^*} ((C\Pi)_w)^T \Pi (C\Pi)_w \right)^* \leq \Pi \\ &\iff \forall w \in \mathcal{L}^* \quad ((C\Pi)_w)^T \Pi (C\Pi)_w \leq \Pi \\ &\iff \forall \alpha \in \mathcal{L} \quad \Pi C_\alpha^T \Pi C_\alpha \leq \Pi \\ &\iff \forall \alpha \in \mathcal{L} \quad C_\alpha^T \Pi C_\alpha \leq \Pi \end{aligned}$$

We now notice that for any $I, J \in \Pi$ and $\alpha \in \mathcal{L}$,

$$\begin{aligned} \alpha L_I \cap L_{V^2} \subseteq L_J &\iff \\ &\iff \left(\alpha \bigcup_{\langle x, y \rangle \in I} L(x, y) \right) \cap L_{V^2} \subseteq L_J \\ &\iff \forall \langle \bar{x}, \bar{y} \rangle \in I \quad \alpha L(\bar{x}, \bar{y}) \cap L_{V^2} \subseteq L_J \\ &\iff \forall \langle \bar{x}, \bar{y} \rangle \in I \quad \forall \langle x, y \rangle \in V^2 \quad \alpha L(\bar{x}, \bar{y}) \cap L(x, y) \subseteq L_J \\ &\iff \forall \langle \bar{x}, \bar{y} \rangle \in I \quad \forall \langle x, y \rangle \in V^2 \quad \alpha L(\bar{x}, \bar{y}) \cap L(x, y) \neq \emptyset \implies \langle x, y \rangle \in J, \end{aligned}$$

and the last condition is straightforwardly equivalent to $C_\alpha^T \Pi C_\alpha \leq \Pi$. ■

We are now ready to characterize graphs which admit monodrome, left regular partitions coarser than T .

Definition 4 Let U be the least partition containing T and closed under $(-)^{C_{\mathcal{L}}}$, i.e., U is the least partition containing T such that $U^{C_{\mathcal{L}}} \leq U$.

Note that since $U^{C_{\mathcal{L}}} \geq U$ by definition, closure under $(-)^{C_{\mathcal{L}}}$ can be equivalently stated by saying that $U^{C_{\mathcal{L}}} = U$. U certainly exists because the class of partitions coarser than T and closed under $(-)^{C_{\mathcal{L}}}$ is nonempty and closed under intersection.

Theorem 5 A graph G with colouring λ has sense of direction iff U is monodrome.

Proof. If G has sense of direction then by Theorem 4 there is a left regular monodrome partition coarser than T and, by minimality, coarser than U ; thus, U is monodrome. For the other direction, just take $\Pi = U$ and use again Theorem 4. ■

Finally, we show how to explicitly compute U :

Theorem 6 $U = (T)^{(C_{\mathcal{L}})^{n^2}}$.

Note that we used an exponential notation for the iteration of the operator $(-)^{C_{\mathcal{L}}}$: i.e., $(T)^{(C_{\mathcal{L}})^{n^2}} = (\dots((T)^{C_{\mathcal{L}}})^{C_{\mathcal{L}}}\dots)^{C_{\mathcal{L}}}$ (n^2 times).

Proof. By a trivial argument of domain theory, U can be obtained by iterating the $(-)^{C_{\mathcal{L}}}$ operator over T until a fixed point is reached. However, T can have at most n^2 equivalence classes, and each application of $(-)^{C_{\mathcal{L}}}$ decreases this number by at least 1, hence the thesis. ■

7 The decision problem for sense of direction

Using the results of the previous section we shall now establish that one can decide in (sequential) polynomial time whether a given coloured graph G with n nodes has sense of direction.

7.1 Checking left regularity

There are two main steps for computing $(-)^{C_{\mathcal{L}}}$: determining the matrices (relations) C_{α} , and computing $\sum_{\alpha \in \mathcal{L}} C_{\alpha} \otimes C_{\alpha}$ (with \otimes we denote the Kronecker product of matrices). This because of the following

Proposition 3 Given a relation $R \subseteq V^2 \times V^2$, let \hat{R} be defined by

$$\hat{R} = \left(\sum_{\alpha \in \mathcal{L}} C_{\alpha} \otimes C_{\alpha} \right) (R \otimes R).$$

Then, we have that

$$\langle x, y \rangle \sum_{w \in \mathcal{L}^*} ((CR)_w)^T R (CR)_w \langle x', y' \rangle \quad (1)$$

iff there are $\langle \bar{x}, \bar{y} \rangle, \langle \bar{x}', \bar{y}' \rangle$ such that $\langle \bar{x}, \bar{y} \rangle R \langle \bar{x}', \bar{y}' \rangle$ and

$$\langle \langle \bar{x}, \bar{y} \rangle, \langle \bar{x}', \bar{y}' \rangle \rangle \hat{R}^* \langle \langle x, y \rangle, \langle x', y' \rangle \rangle \quad (2)$$

Proof. Since

$$\hat{R} = \left(\sum_{\alpha \in \mathcal{L}} C_\alpha \otimes C_\alpha \right) (R \otimes R) = \sum_{\alpha \in \mathcal{L}} C_\alpha R \otimes C_\alpha R,$$

we have that (2) holds iff there is a word $w = \alpha_1 \alpha_2 \cdots \alpha_n \in \mathcal{L}^*$ such that

$$\langle \bar{x}, \bar{y} \rangle, \langle \bar{x}', \bar{y}' \rangle \rangle C_{\alpha_1} R C_{\alpha_2} R \cdots C_{\alpha_n} R \otimes C_{\alpha_1} R C_{\alpha_2} R \cdots C_{\alpha_n} R \langle \langle x, y \rangle, \langle x', y' \rangle \rangle,$$

i.e., such that

$$\langle \bar{x}, \bar{y} \rangle (CR)_w \langle x, y \rangle \wedge \langle \bar{x}', \bar{y}' \rangle (CR)_w \langle x', y' \rangle. \quad (3)$$

But (1) is equivalent to the existence of a word $w \in \mathcal{L}^*$ such that $\langle x, y \rangle ((CR)_w)^T R (CR)_w \langle x', y' \rangle$, which in turn is equivalent to the existence of pairs $\langle \bar{x}, \bar{y} \rangle R \langle \bar{x}', \bar{y}' \rangle$ satisfying (3), or equivalently (2). ■

This means that, given the matrices C_α , we can compute $R^{C_{\mathcal{L}}}$ as follows:

- compute $(R \otimes R) \hat{R}^*$;
- create a new $n^2 \times n^2$ matrix S such that $S(\langle x, y \rangle, \langle x', y' \rangle) = 1$ iff there is a $\langle \bar{x}, \bar{y} \rangle \in V^2$ such that $\langle \langle \bar{x}, \bar{y} \rangle, \langle \bar{x}, \bar{y} \rangle \rangle (R \otimes R) \hat{R}^* \langle \langle x, y \rangle, \langle x', y' \rangle \rangle$;
- compute S^* , which is exactly $R^{C_{\mathcal{L}}}$.

We start by showing how to compute the C_α 's by means of the same techniques we used in proving Proposition 1. We shall add to each node of G a bunch of $|\mathcal{L}|$ incoming arcs starting from new “artificial” nodes, and apply again the product construction.

Proposition 4 Let the graph $H = (W, B)$ be defined as follows: $W = V + V \times \mathcal{L}$, $B = A \cup \{ \langle \langle x, \alpha \rangle, x \rangle \mid x \in V, \alpha \in \mathcal{L} \}$. Let us extend the colouring λ of G to a colouring $\mu : B \rightarrow \mathcal{L}$ by setting $\mu(\langle \langle x, \alpha \rangle, x \rangle) = \alpha$. Let D be the $(n + n|\mathcal{L}|)^2 \times (n + n|\mathcal{L}|)^2$ matrix

$$D(\langle \langle s, s' \rangle, \langle t, t' \rangle \rangle) = 1 \iff \langle s, t \rangle, \langle s', t' \rangle \in B \wedge \mu(\langle s, t \rangle) = \mu(\langle s', t' \rangle)$$

where s, s', t, t' range over W . Then, for all $x, x', y, y' \in V$

$$D^*(\langle \langle x, \alpha \rangle, x' \rangle, \langle y, y' \rangle \rangle) = 1 \iff \alpha L(x, y) \cap L(x', y') \neq \emptyset.$$

Proof. We proceed like in the proof of Proposition 1. $D^*(\langle \langle x, \alpha \rangle, x' \rangle, \langle y, y' \rangle \rangle) = 1$ happens iff there is a word $w = \alpha_1 \alpha_2 \cdots \alpha_n$ ($n > 0$) such that $\langle \langle x, \alpha \rangle, x' \rangle \cdot w = y$ and $x' \cdot w = y'$. But again this can happen iff $\alpha_1 = \alpha$, and since $\langle x, \alpha \rangle \cdot \alpha = x$, we have that $\alpha_2 \cdots \alpha_n \in L(x, y)$. But then $\alpha L(x, y) \ni w \in L(x', y')$. ■

Corollary 1 $C_\alpha(\langle x, y \rangle, \langle x', y' \rangle) = D^*(\langle \langle x, \alpha \rangle, x' \rangle, \langle y, y' \rangle \rangle)$.

We now have the following

Theorem 7 Given a relation $R \subseteq V^2 \times V^2$, $R^{C_{\mathcal{L}}}$ can be computed in logarithmic time using a polynomial number of processors.

Proof. Recalling that $|\mathcal{L}|$ is polynomially bounded in n , it is easy to see that the matrix D (with the notation of Proposition 4) can be built in constant time using a polynomial number of processors. Then, we just have to perform a series of transitive closures, products and permutations, which can be computed in logarithmic time using a polynomial number of processors. The details are left to the reader. ■

7.2 Sense of direction is in P

We now come to the main theorems of this section. We start again by formalizing our problem:

Problem 2 SENSE OF DIRECTION

Instance: A graph G with colouring λ .

Question: Is λ a sense of direction?

Theorem 8 SENSE OF DIRECTION is in P.

Proof. We can compute T in logarithmic time using n^6 processors; thus, we can *a fortiori* compute it sequentially in polynomial time. We just have to iterate n^2 times $(-)^{C_{\mathcal{L}}}$ over T in order to compute U (see Definition 4), and this can be done again in polynomial time by Theorem 7. By Theorem 5, we can decide whether our graph has sense of direction by checking the monodromy of U , which can be done in $O(n^3)$ steps. ■

Note that since the operator $(-)^{C_{\mathcal{L}}}$ is computable in parallel logarithmic time, and the bound on the number of iterations provided by Theorem 6 is very rough, it is an interesting open problem to characterize large classes of graphs for which a (poly)logarithmic number of iterations suffices.

8 Cayley graphs and open problems

The positive results we presented about the decision problem for (weak) sense of direction of a *coloured* graph should not obscure the fact that the next and most important open problem is to decide how many colours are really necessary in order to give (weak) sense of direction to a given graph. We present some partial results along this line, which highlight a surprising connection with the open problem of recognizing Cayley graphs, and provide a fast parallel algorithm for recognizing Cayley colour graphs. Our problem is as follows:

Problem 3 k -(WEAK) SENSE OF DIRECTION

Instance: A graph G and a positive integer k .

Question: Is there a (weak) sense of direction λ for G using at most k colours?

The number of colours which are necessary in order to give to a graph G with n vertices a (weak) sense of direction is always bounded from below by the maximum outdegree of the graph, and from above by the number of vertices (order arbitrarily the vertices and then assign to each arc from the i -th to the j -th vertex the colour $(i - j) \bmod n$). When a graph G possesses a (weak) sense of direction λ achieving the lower bound, we say that λ is *minimal*. Clearly the following problem is polynomial time reducible to k -(WEAK) SENSE OF DIRECTION:

Problem 4 MINIMAL (WEAK) SENSE OF DIRECTION**Instance:** A graph G .**Question:** Is there a minimal (weak) sense of direction λ for G ?

Let now Γ be a finite group, and $S \subseteq \Gamma$ a set of elements of Γ (which are usually taken to be generators). The *Cayley graph* of Γ with respect to S is the graph having vertex set Γ , and an arc from g to h whenever there is an $s \in S$ such that $gs = h$. When S is closed under inversion, the resulting graph is symmetric (or, equivalently, it is undirected). Note that the graph is $|S|$ -regular (i.e., there are exactly $|S|$ arcs going into and $|S|$ arcs going out of every vertex), and that it is strongly connected iff S is a set of generators.

Each Cayley graph has a natural deterministic colouring in the set S : an arc from g to gs is coloured by s . Such a graph is called the *Cayley colour graph* of Γ with respect to S , and it has minimal sense of direction [Flo96]. Consider now the problem of recognizing Cayley colour graphs:

Problem 5 CAYLEY COLOUR GRAPH**Instance:** A graph G with colouring λ .**Question:** Is G a Cayley colour graph?

In [BV97] the following result is proved:

Theorem 9 A coloured graph is a Cayley colour graph iff it is outregular and has minimal (weak) sense of direction.

As noted in [BV97], since non-outregular graphs and graphs using more colours than the outdegree can be detected in AC^1 , we can use the algorithm described in Section 5 in order to recognize Cayley colour graphs. This implies that

Theorem 10 CAYLEY COLOUR GRAPH is in AC^1 .

However, there is something more: Theorem 9 implies that a graph is a Cayley graph iff it is outregular and it can be given a minimal (weak) sense of direction. Thus, the following problem is equivalent to MINIMAL (WEAK) SENSE OF DIRECTION, and therefore reducible to k -(WEAK) SENSE OF DIRECTION:

Problem 6 CAYLEY GRAPH**Instance:** A graph G .**Question:** Is G a Cayley graph?

We currently conjecture that CAYLEY GRAPH is at least as hard as graph isomorphism. To our knowledge, Theorem 10 is the first complexity result on Cayley colour graphs, and (surprisingly!) no results are known about CAYLEY GRAPH.

9 Acknowledgments

We would like to thank Carlo “Mere” Mereghetti for many useful discussions and Gianfrancesco “Frank” Ravasio for his careful proofreading.

References

- [AvLSZ89] Hagit Attiya, Jan van Leeuwen, Nicola Santoro, and Shmuel Zaks. Efficient elections in chordal ring networks. *Algorithmica*, 4:437–446, 1989.
- [BV97] Paolo Boldi and Sebastiano Vigna. Minimal sense of direction and decision problems for Cayley graphs. *Inform. Process. Lett.*, 64(6):299–303, 1997.
- [Flo96] Paola Flocchini. Minimal sense of direction in regular graphs. Submitted, 1996.
- [FMS95] Paola Flocchini, Bernard Mans, and Nicola Santoro. Sense of direction: Formal definition and properties. In Paola Flocchini, Bernard Mans, and Nicola Santoro, editors, *Structure, Information and Communication Complexity. Proc. 1st Colloquium SIROCCO '94*, pages 9–33. Carleton University Press, 1995.
- [KR90] R. Karp and V. Ramachandran. A survey of parallel algorithms for shared-memory machines. In J. Van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, pages 869–941. North–Holland, 1990.
- [LMW86] M.C. Loui, T.A. Matsushita, and D.B. West. Election in complete networks with a sense of direction. *Information Processing Letters*, 22:185–187, 1986. See also *Information Processing Letters*, vol.28, p.327, 1988.
- [San84] Nicola Santoro. Sense of direction, topological awareness and communication complexity. *SIGACT News*, 2(16):50–56, 1984.
- [Tar72] Robert E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Comput.*, 1(2):146–160, 1972.