Combinatorics, Probability and Computing (2004) 13, 123–136. © 2004 Cambridge University Press
DOI: 10.1017/S0963548303006011 Printed in the United Kingdom

Distributed Colouring and Communication in Rings with Local Knowledge

ANDERS DESSMARK[1†] and ANDRZEJ PELC[2‡§]

[1] Department of Computer Science, Lund Institute of Technology, Box 118, S-22100 Lund, Sweden
(e-mail: andersd@cs.lth.se)

[2] Département d'informatique, Université du Québec en Outaouais, Hull, Québec J8X 3X7, Canada
(e-mail: pelc@uqo.ca)

Received 3 June 2002; revised 10 April 2003

We consider two interrelated tasks in a synchronous n-node ring: *distributed constant colouring* and *local communication*. We investigate the impact of the amount of knowledge available to nodes on the time of completing these tasks. Every node knows the labels of nodes up to a distance r from it, called the *knowledge radius*. In distributed constant colouring every node has to assign itself one out of a constant number of colours, so that adjacent nodes get different colours. In local communication every node has to communicate a message to both of its neighbours. We study these problems in two popular communication models: the *one-way* model, in which each node can only either transmit to one neighbour or receive from one neighbour, in any round, and the *radio* model, in which simultaneous receiving from two neighbours results in interference noise. Hence the main problem in fast execution of the above tasks is breaking symmetry with restricted knowledge of the ring.

We show that distributed constant colouring and local communication are tightly related and one can be used to accomplish the other. Also, in most situations the optimal time is the same for both of them, and it strongly depends on the knowledge radius. For knowledge radius $r = 0$, *i.e.*, when each node knows only its own label, our bounds on time for both tasks are tight in both models: the optimal time in the one-way model is $\Theta(n)$, while in the radio model it is $\Theta(\log n)$. For knowledge radius $r = 1$ both tasks can be accomplished in time $O(\log\log n)$ in the one-way model, if the ring is oriented. For $2 \leqslant r \leqslant c \log^* n$, where $c < 1/2$, the upper bounds on time are $O(\log^{(2r)} n)$ in the one-way model and $O(\log^{(2\lfloor r/2 \rfloor)} n)$ in the radio model; the lower bound is $\Omega(\log^* n)$, in both models. For $r \geqslant (\log^* n)/2$ both tasks can be completed in constant time, in the one-way model, and distributed constant colouring also in the radio model. Finally, if $r \geqslant \log^* n$ then constant time is also enough for local communication in the radio model.

† This research was done during the stay of Anders Dessmark at the Université du Québec en Outaouais as a postdoctoral fellow.
‡ Supported in part by NSERC grant OGP 0008136.
§ A preliminary version of this paper appeared in *Proc. International Parallel and Distributed Processing Symposium* (IPDPS 2001), San Francisco, USA, April 2001.

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

1. Introduction

1.1. The background

Coping with incomplete information in the design of distributed algorithms and in network communication is an important problem whose various aspects have been extensively studied. The most extreme case arises when nodes are completely ignorant of network topology and do not know even their immediate neighbourhood. Such a scenario was considered, *e.g.*, in the context of graph exploration [1, 25] and of broadcasting in radio networks [8, 9]. More often, and more realistically, partial knowledge of the network, usually of local nature, is assumed. In [6] the authors studied time of broadcasting in radio networks whose nodes know only their immediate neighbourhood, while in [5] the aim was to investigate trade-offs between the size of the part of the network known to nodes, and message complexity of broadcasting. In [20, 21] local knowledge available to nodes was used to distributedly solve problems such as fault mending and graph colouring.

1.2. The problem and previous work

In this paper we study the impact of the amount of local knowledge available to nodes on the execution time of distributed tasks of local nature. (This should be contrasted, *e.g.*, with [5] where the number of messages was selected as the efficiency measure.) In order to focus exclusively on trade-offs between knowledge and speed of execution of these tasks (as opposed to difficulties caused by the structure of the network), we work in one of the simplest topologies, the ring, as, *e.g.*, in [3, 11] for anonymous networks.

The primary object of our study is the task of *node-colouring*, that is, finding an assignment of colours to nodes so that adjacent nodes get different colours. While finding a colouring of the ring with 3 colours is straightforward if the algorithm is centralized, or equivalently, if nodes have global knowledge of the ring, this task becomes much more difficult if colouring has to be found in a distributed way, and nodes have *a priori* only local knowledge. We investigate distributed node-colouring with a constant number of colours (called *distributed constant colouring* throughout the paper), assuming only local knowledge available to nodes.

Node-colouring is a fundamental graph problem, and hence distributed algorithms for this task have been extensively studied [4, 10, 17, 21, 22, 23, 24]. In all these papers a very strong computational model has been used. We will call it *Linial's model* (*cf.* the description in [21]). This is a synchronous model whose main feature is that, in every round, every node can execute an arbitrary amount of local computations and can send and receive messages to and from all of its neighbours. Hence, although *a priori* every node knows only its own label, it can get knowledge of the part of the graph at distance r from it, in time r. As mentioned in [21], this model is of purely theoretical interest. In more realistic models communication is usually slower and hence acquiring knowledge of the graph is more time-consuming.

Cole and Vishkin [10] described an algorithm for the PRAM model achieving a 3-colouring of an n-node ring in time $O(\log^* n)$. This algorithm is valid in Linial's model as well. Linial [21] showed a matching lower bound $\Omega(\log^* n)$ on time needed for 3-colouring

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

of the ring. In the same paper he showed a deterministic $O(\Delta^2)$-colouring in time $O(\log^* n)$, for arbitrary n-node graphs of maximum degree Δ. The latter result was previously shown in [17] only for graphs of constant degree.

Both randomized and deterministic distributed node-colouring have been also studied, *e.g.*, in [4, 24].

A closely related problem is that of *colour reduction*. Given a c-colouring, obtain a $f(c)$-colouring in 1 round, for $f(c) < c$. The solution of this problem implies $f(n)$-colouring in 1 round, since labels can be considered as n different colours. Cole and Vishkin [10] and Goldberg, Plotkin and Shannon [17] showed a method that achieves a colour reduction from c colours to $O((\log c)^\Delta)$ colours. Linial's method [21] reduces the number of colours from c to $O(\Delta^2 \log c)$. Mayer, Naor and Stockmeyer [23] showed a (constructive) colour reduction from c to $O(\Delta 2^{3\Delta} \log\log c)$ colours.

We work in communication models much more restrictive than Linial's model used in the above papers (see Section 1.3). Hence the time of colouring strongly depends on the time required to communicate with neighbours (which can be done in a single round in Linial's model). This basic communication task consisting of exchanging information between neighbours will be called *local communication*. (It is also sometimes called *neighbourhood gossiping* [15] or *single round simulation* [2].) It is important on its own, since its repetitive execution enables each node to gain knowledge about the part of the network at distance t from it. In particular, fast local communication makes possible fast colouring. It turns out that – in restrictive communication models – the dependence is reciprocal: the availability of node-colouring with few colours enables fast local communication. Thus we study both tasks simultaneously. In some of our algorithms both tasks are interleaved: an obtained colouring permits faster communication, which in turn makes possible improvement of the previous colouring.

We consider only deterministic algorithms. Results concerning randomized colouring have been mentioned before, and randomized communication in networks in various models has been considered, *e.g.*, in [2, 6, 13]. Randomized communication algorithms usually work without assuming any knowledge of the network.

1.3. The models

We assume that communication is synchronous: nodes transmit messages in rounds controlled by a global clock and a message sent in each round is available at the neighbour in the next round. Since our aim is to investigate execution time for distributed tasks, it should be clear that efficiency of algorithms will depend on the communication model, which specifies what actions can be executed by nodes in a single round. Many different models have been used in the literature and are described in surveys [14, 18]. In the *all-port* model, for example, in which each node can transmit to all neighbours and receive from all neighbours simultaneously (*cf.* Linial's model described in Section 1.2), the task of exchanging messages between neighbours can be trivially completed in one round (without any knowledge of the network). Local communication becomes harder in more restrictive models in which a node can hear a message from its neighbour only if no other message is sent to it in the given round. Models satisfying this assumption can be

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

divided into two classes: models for point-to-point networks, in which a node can choose on which of the adjacent links it transmits in a given round, and radio models in which a message sent by a node reaches all of its neighbours (if no collision occurs). For our study we select one model from each class. Below we describe these models and justify the reasons for our choice.

The *one-way* model (also called *one-port half-duplex*) is one of the most popular models for point-to-point networks in the literature on broadcasting and gossiping, studied, *e.g.*, in [12, 19]. In every round each node can only either transmit a message to one neighbour or receive a message from one neighbour but not both. Any attempt to violate this rule results in a run-time error, *i.e.*, a correct communication algorithm in the one-way model must never produce such a collision. The one-way model has the advantage of being the weakest of all models for point-to-point networks (*cf.* [14, 18]) hence any algorithm designed for it works for all other models of this class, with at least the same efficiency.

The *radio* model is defined as follows. In every round every node acts either as a *transmitter* or as a *receiver*. A node acting as a transmitter sends a message to all of its neighbours. A node acting as a receiver in a given round gets a message if and only if exactly one of its neighbours transmits in this round. If at least two neighbours v and v' of u transmit in a given round, none of the messages is received by u in this round. In this case we say that a *collision* occurred at u. The radio model was studied, *e.g.*, in [6, 7, 8] in the context of broadcasting. Local communication in the radio model was considered in [2] for arbitrary graphs, under the name *single round simulation*. Two variations of the radio model are possible. In the case of a collision at u, node u may either hear nothing (except for the background noise), or it may receive *interference noise* different from any message received properly but also different from background noise. These two variations are often referred to as the absence (resp. availability) of *collision detection* (*cf.*, *e.g.*, [6]). A discussion justifying both variations can be found in [6, 16]. We use the radio model without collision detection (as, *e.g.*, in [6]) but all our results hold in the other variation as well.

It remains to describe the measure of local knowledge available to nodes. We assume that all nodes of an n-node ring have distinct labels which are integers between 1 and n. We identify labels with their binary representations. Our algorithms and arguments are easy to modify if labels are in the range 1 to M, where $M \in O(n)$. (It is easy to see that if the ring is anonymous, local communication and distributed colouring are impossible in any of the above models.) Suppose that labels of consecutive nodes are $l_0, \ldots, l_{n-1}$. For a natural number r we say that r is the *knowledge radius* of the ring if every node v with label l_i knows the sequence of labels $l_{i-r}, \ldots, l_i, \ldots, l_{i+r}$, where arithmetic on indices is modulo n. Thus radius 0 corresponds to every node knowing only its own label, and radius $\lfloor n/2 \rfloor$ corresponds to full knowledge of the ring. In the case of positive knowledge radius there is a further distinction between availability and lack of *orientation*. Apart from knowing the sequence $l_{i-r}, \ldots, l_i, \ldots, l_{i+r}$ node v may know which of the nodes with labels l_{i-1} l_{i+1} is its *left* and which is its *right* neighbour, or else nodes may have no such knowledge. (In the first case we assume that left/right notions are coherent among nodes, *i.e.*, the left neighbour of the right neighbour of v is v itself.) All our results except one

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

Table 1. Summary of results: optimal time of distributed constant colouring and of local communication, depending on knowledge radius

Knowledge radius	Model	Time
$r = 0$	one-way	$\Theta(n)$
	radio	$\Theta(\log n)$
$r = 1$	one-way	$O(\log \log n)$ (oriented ring)
$2 \leqslant r \leqslant c \log^* n$ ($c < 1/2$)	one-way	$O(\log^{(2r)} n)$ $\Omega(\log^* n)$
	radio	$O(\log^{(2\lfloor r/2 \rfloor)} n)$ $\Omega(\log^* n)$
$r \geqslant (\log^* n)/2$	one-way	$O(1)$
	radio	distributed constant colouring: $O(1)$
$r \geqslant \log^* n$	radio	local communication: $O(1)$

are the same for the oriented and unoriented case, and thus we present algorithms for unoriented rings.

1.4. Overview of results and open problems

We establish upper and lower bounds on the execution time of two interrelated distributed tasks performed in an n-node ring: distributed colouring with a constant number of colours, and local communication. These bounds strongly depend on knowledge radius. For each upper bound we give an algorithm for each of these tasks, working in the given time. In what follows, $\log^{(i)} n$ denotes the ith iteration of the logarithm function, thus, *e.g.*, $\log^{(2)} n = \log \log n$.

For knowledge radius $r = 0$, *i.e.*, when each node knows only its own label, our bounds on time for both tasks are tight in both models: the optimal time in the one-way model is $\Theta(n)$, while in the radio model it is $\Theta(\log n)$. For knowledge radius $r = 1$ both tasks can be accomplished in time $O(\log \log n)$ in the one-way model, if the ring is oriented. For $2 \leqslant r \leqslant c \log^* n$, where $c < 1/2$, the upper bounds on time are $O(\log^{(2r)} n)$ in the one-way model and $O(\log^{(2\lfloor r/2 \rfloor)} n)$ in the radio model; the lower bound is $\Omega(\log^* n)$, in both models. For $r \geqslant (\log^* n)/2$ both tasks can be completed in constant time, in the one-way model, and distributed constant colouring also in the radio model. Finally, if $r \geqslant \log^* n$ then constant time is also enough for local communication in the radio model.

Table 1 summarizes our results. It also indicates problems that remain open. While the algorithms for knowledge radius $r > 1$ are very fast, some gaps remain between upper and lower bounds in this case. In particular, we do not know if any of the considered tasks can be completed in time $O(\log^* n)$ for constant knowledge radius. Another challenging open problem is determining whether any of our tasks can be completed in the unoriented ring in sublinear time for knowledge radius 1 in the one-way model. We conjecture that this is impossible, as opposed to the oriented case.

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

2. Positive knowledge radius

In this section we establish upper and lower bounds on optimal time of distributed constant colouring and of local communication for positive knowledge radius, depending on its size. Our algorithms are based on the following idea. Using the given knowledge radius r, obtain (without need of communication) a colouring with the number of colours possible to obtain in Linial's model in r rounds. Then use this colouring to obtain a fast local communication procedure. Next apply this procedure several times to acquire a larger knowledge radius, which in turn yields a better colouring, and so on. Thus colouring and communication are improved iteratively, serving both as a target and as a tool.

The following lemmas on the relationship between colouring and local communication will be helpful throughout the section. Their proofs are left to the reader.

Lemma 2.1. *Assume knowledge radius $r \geqslant 2$ in the one-way model. If a c-colouring can be done in time t then local communication can be performed in time $t + 4c$.*

In the radio model, the inability to choose direction of transmissions makes the straightforward colouring of the ring insufficient for avoiding collisions. We will instead use a colouring of the ring augmented with edges between every pair of nodes at distance 2. Call this graph the *augmented ring*. In other words, we consider a colouring of nodes of the (normal) ring, assigning different colours to nodes at distance $\leqslant 2$. □

Lemma 2.2. *Assume knowledge radius $r \geqslant 2$ in the radio model. If nodes of the ring can be coloured with c colours in time t, in such a way that all nodes within distance 2 have different colours, then local communication in the ring can be performed in time $t + c$.*

We first show that, for large knowledge radius, both tasks can be accomplished in constant time, in both models. □

Proposition 2.3. *Assume knowledge radius $r \geqslant \log^* n$ in the radio model. Local communication in an n-node ring can be completed in constant time.*

Proof. It follows from [23] that colour reduction from c to $O(\log \log c)$ colours can be done constructively (in one round) in Linial's model, for all bounded degree graphs. This applies to the augmented ring, as it has degree 4. Knowledge radius $r \geqslant \log^* n$ in the ring implies knowledge radius $r' \geqslant \lfloor (\log^* n)/2 \rfloor$ in the augmented ring. Hence, iterating this colour reduction $\lfloor (\log^* n)/2 \rfloor$ times brings down the number of colours from n to constant, with only internal computations, and hence in constant time, also in the radio model. (Indeed, during $\lfloor (\log^* n)/2 \rfloor$ rounds performed in Linial's model, a node can only get information on nodes at distance $\lfloor (\log^* n)/2 \rfloor$. Knowledge radius $\lfloor (\log^* n)/2 \rfloor$ guarantees this information *a priori*, and hence no communication is needed to simulate these $\lfloor (\log^* n)/2 \rfloor$ rounds.) This final colouring, together with Lemma 2.2, implies the result. □

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

To obtain distributed constant colouring in both models, and local communication in the one-way model, we do not need to consider the augmented ring, and hence knowledge radius twice shorter is sufficient. The following result is a consequence of $\lfloor(\log^* n)/2\rfloor$ iterations of colour reduction from [23], applied to the ring.

Proposition 2.4. *Assume knowledge radius $r \geqslant (\log^* n)/2$. Distributed constant colouring of an n-node ring can be done in constant time in both the one-way and the radio models.*

Proposition 2.4 and Lemma 2.1 imply the following. □

Proposition 2.5. *Assume knowledge radius $r \geqslant (\log^* n)/2$ in the one-way model. Local communication in an n-node ring can be completed in constant time.*

For knowledge radius smaller than above but larger than 1, the iterated colour reduction used previously accomplishes both tasks in time depending on r. □

Proposition 2.6. *Assume knowledge radius $r \geqslant 2$ in the one-way model. Distributed constant colouring and local communication can be completed in time $O(\log^{(2r)} n)$ in an n-node ring.*

Proof. The proof is by induction on $i = \lfloor(\log^* n)/2\rfloor - r$. The proposition is true for all $r \geqslant (\log^* n)/2$ by Propositions 2.4 and 2.5. Assume that the theorem is true for all $r > r' \geqslant 2$. By r' iterations of the colour reduction procedure from [23], we obtain an $O(\log^{(2r')} n)$-colouring (without any communication). We now apply Lemma 2.1 and perform local communication in time $O(\log^{(2r')} n)$. During this local communication, every node v informs its neighbours of the labels of nodes at distance r' from v. Thus, provided with this additional information, nodes now have knowledge radius $r' + 1$. By the induction hypothesis, distributed constant colouring can now be performed in time $O(\log^{(2(r'+1))} n)$, and hence the total time is $O(\log^{(2r')} n)$, which completes the proof. □

For the radio model we need a radius twice as large to get the same time, as before. The proof is similar to that of Proposition 2.6. It is left to the reader.

Proposition 2.7. *Assume knowledge radius $r \geqslant 2$ in the radio model. Distributed constant colouring and local communication in an n-node ring can be completed in time $O(\log^{(2\lfloor r/2\rfloor)} n)$.*

We will now show a lower bound on the execution time of both tasks. It is valid for knowledge radius $r \leqslant c \log^* n$, where $c < 1/2$. □

Theorem 2.8. *Every distributed constant colouring algorithm and every local communication algorithm, working in an n-node ring, for knowledge radius $r \leqslant c \log^* n$, where $c < 1/2$, have running time $\Omega(\log^* n)$, in both models of communication.*

Proof. We first prove the result for distributed constant colouring. Let c' be a positive constant smaller than $1/2 - c$. Consider either the one-way or the radio model. Assume

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

that distributed constant colouring can be done in time smaller than $c' \log^* n$. Hence it can be done in Linial's model, in time $d \log^* n$, where $d = c + c' < 1/2$. Let g be the number of colours. Hence, in time $d \log^* n + d'$, for some constant d', we can get a 3-colouring of the ring in Linial's model, by iterating the colour reduction procedure from [23] d' times (to bring the number of colours from g to 3). For sufficiently large n, we have $d \log^* n + d' < 1/2(\log^* n - 3)$. This gives a 3-colouring of the ring in time less than $1/2(\log^* n - 3)$ in Linial's model, contradicting the lower bound from [21].

We now consider local communication. Let c'' be a positive constant smaller than $(1/2 - c)/3$. We give a common argument for both the one-way and the radio models, although the proof could be slightly simplified for the one-way model. Assume that local communication can be done in time smaller than $c'' \log^* n$ by algorithm A. We will show a procedure which obtains a colouring with few colours fast, using this local communication algorithm A. This will yield a contradiction with the lower bound from [21], similar to the above.

Run algorithm A. For any node v with neighbours x and y, let t_x^v (respectively t_y^v) be the round numbers in which v first gets a message from x (respectively from y). Run A the second time. In this run, every node v informs both neighbours x and y of t_x^v and t_y^v. After this second run, every node v knows t_v^x and t_v^y. (Note the change of superscript and subscript.) Let $t_v = \min(t_v^x, t_v^y)$. The integer t_v must be different either from t_x or from t_y. (Otherwise, x and y would transmit in the same round in which at least one of them got the message from v.) Assume without loss of generality that $t_v \neq t_x$. Run A the third time, informing both neighbours of v of the integer t_v. Now each node v knows t_v, t_x, and t_y. Node v chooses the pair $(t_v, 0)$ as its colour, unless $t_v = t_y$ and the label of v is larger than that of y, in which case v chooses the pair $(t_v, 1)$ as its colour. Hence neighbours get different colours, and thus we obtain an h-colouring of the ring, for $h = 2c'' \log^* n$, in time $3c'' \log^* n$.

Hence an h-colouring can be obtained in Linial's model in time at most $r + 3c'' \log^* n \leqslant f \log^* n$, where $f \leqslant c + 3c'' < 1/2$. Let f' be the minimum number of iterations of the colour reduction procedure from [23], needed to bring the number of colours from h to 3. Hence a 3-colouring of the ring can be obtained in time $f \log^* n + f'$ in Linial's model. Since $h < \log^* n$, we have $f' \leqslant a \log^*(\log^* n)$, for some constant a. For sufficiently large n, $f \log^* n + f' < 1/2(\log^* n - 3)$. As before, this contradicts the lower bound from [21]. □

We conclude this section by observing an upper bound on the time of distributed constant colouring and of local communication for knowledge radius 1, *i.e.*, when nodes know labels of their neighbours.

Proposition 2.9. *Distributed constant colouring and local communication can be completed in time $O(\log \log n)$, in an n-node oriented ring with knowledge radius 1, in the one-way model.*

Proof. The colour reduction described in [23] gives the colouring of the ring with $O(\log \log n)$ colours in 1 round in Linial's model. Hence, with knowledge radius 1, this colouring can be obtained without communication (also in the one-way model). This

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

implies that the following local communication procedure works in time $O(\log \log n)$. It proceeds in $k \in O(\log \log n)$ phases, where k is the number of colours in this colouring. In the ith phase, nodes with colour i transmit first left then right. After one application of this procedure, knowledge radius 2 is acquired. Then, in view of Proposition 2.6, distributed constant colouring can be accomplished in time $O(\log^{(4)} n)$. Thus the total time for distributed constant colouring is also $O(\log \log n)$. □

If the ring is not oriented, the above described procedure cannot be applied. In this case, the gaps between the upper bound $O(n)$ and lower bound $\Omega(\log^* n)$ are very large. We conjecture that both our tasks require linear time in the unoriented ring with knowledge radius 1, for the one-way model.

3. Knowledge radius 0

In this section we consider the tasks of distributed constant colouring and of local communication in the case when each node knows only its own label. This scenario differs significantly from the case when knowledge radius is positive. Indeed, for $r > 0$, it was possible to use the results concerning colouring in Linial's model, since the only effect of the strong communication power of this model can be acquiring knowledge radius r in r rounds. However, for knowledge radius 0 in our restrictive models, even the results concerning colouring in 1 round in Linial's model cannot be used immediately because nodes do not have *a priori* knowledge which can be obtained in Linial's model in this single round of communication. Thus the algorithms must proceed by first independently establishing local communication and then using it to get distributed constant colouring. It turns out that this difficulty considerably slows down both tasks, as compared to the scenario with positive knowledge radius.

Our first result gives tight bounds on optimal time of distributed constant colouring and of local communication in the one-way model.

Theorem 3.1. *Optimal time of distributed constant colouring and of local communication in an n-node ring with knowledge radius $r = 0$ is $\Theta(n)$, in the one-way model.*

Proof. The following simple algorithm performing local communication works in time at most $2n$. There are n phases each consisting of 2 rounds. In the ith phase the node with label i transmits a message to its two neighbours. Since at most one node transmits in each round, there are no collisions.

Distributed constant colouring in time $O(n)$ can be done as follows. First use the above algorithm for local communication twice, thus acquiring knowledge radius 2. Then, by Proposition 2.6, distributed constant colouring can be done in time $O(\log^{(4)} n)$.

In order to prove the lower bound, we first show that every local communication algorithm must use time at least $\lfloor n/3 \rfloor$. Suppose that there exists an algorithm which completes local communication in time $t < \lfloor n/3 \rfloor$. Fix a round $i \leqslant t$ and a node with label v. Fix a numbering 1 and 2 of the two ports of the node. Since we prove the lower bound, the unoriented case follows from the oriented case. Hence assume that the ring

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

is oriented, and port 1 is left while port 2 is right. The knowledge of a node in round i consists of its label v, of the integer i, and of all messages obtained before round i. Thus the algorithm yields a function f specifying the decision made in round i by a node with label v, under the assumption that the node has not received any message before round i. The value $f(i,v)$ of this function can be 0,1, or 2, corresponding to the three decisions: keep silent, transmit on port 1, or transmit on port 2. Call a node with label v an *i-initiator* if $f(i,v) > 0$. Hence an i-initiator is a node that spontaneously transmits in round i.

We show by induction on i ($i \leqslant t$) that the total number of j-initiators, for all $j \leqslant i$, is at most i. First assume $i = 1$. If there were at least two initiators with labels v and w, then it would be possible to position w as a neighbour of v on the side of the port $f(1,v)$. Then in round 1 node with label v transmits to node with label w which also transmits in this round. This yields a collision.

Let S_i be the set of j-initiators, for all $j \leqslant i$. Assume by induction that $|S_i| \leqslant i$, and consider round $i+1 \leqslant t$. It is enough to show that there is at most one $(i+1)$-initiator outside S_i. Suppose that there are at least two such nodes, with labels v and w. Position all nodes from S_i in a segment of the ring. Let a and b be nodes at distance i from respective endpoints of this segment. Consider the segment T of the ring strictly between a and b, not containing S_i. We have $|T| \geqslant n - |S_i| - 2i \geqslant n - 3t \geqslant 2$. If nodes with labels v and w are positioned in T, their distance from any node in S_i exceeds i. Hence none of these nodes could get a message before round $i+1$ and consequently they behave according to the values of function f. Since both these nodes are $(i+1)$-initiators they can be put next to each other in segment T to generate a collision, as for $i = 1$.

It follows that the total number of i-initiators, for all $i \leqslant t$ is at most t. Place all these nodes in a segment of the ring. The same argument as above shows that there exists a node at distance more than t from any of them. This node cannot get any message in the first t rounds, a contradiction. This proves the lower bound $\Omega(n)$ on time of local communication.

We finally prove the lower bound $\Omega(n)$ on time of distributed constant colouring. It is enough to show that – in the oriented ring – any algorithm for distributed constant colouring working in time t yields an algorithm for local communication working in time $t + O(1)$. Indeed, assuming that we have a c-colouring of the ring, the following algorithm performs local communication in time $2c$. The algorithm works in c phases, each consisting of 2 rounds. In phase i, all nodes with colour i transmit first on the right port and then on the left port. In view of the definition of colouring, collisions are avoided. □

We now turn our attention to the radio model. We show that if every node knows only its own label then the optimal time of distributed constant colouring and of local communication is logarithmic in the size of the ring. We first consider the task of local communication and describe an algorithm (working even for unoriented rings) that achieves this performance.

Algorithm Flip-Labels. The algorithm can be divided into two parts. In the first part every node learns the labels of both of its neighbours and learns the message of (at

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

least) one of them. In the second part every node that has not learned messages of both neighbours in the first part, learns the message of the other neighbour.

Part 1 is divided into two phases, each consisting of k steps, where k is the length of the labels. In the ith step of the first phase all nodes that have 0 in the ith position act as transmitters (transmitting their message and their label), and all other nodes act as receivers. In the ith step of the second phase all nodes that have 1 in the ith position act as transmitters, and all other nodes act as receivers. It will be proved that after Part 1 every node v knows the labels of both of its neighbours and the message of at least one neighbour. Call the other neighbour of v its *target*. Every node v compares the labels of both neighbours and computes the first position t where they differ. Call this position the *index* of v.

Part 2 is divided into four stages. In each of them some nodes flip the bit at their index position (from 1 to 0 or from 0 to 1) and then all nodes run Part 1 according to the new labels. Let $I(v)$ be the bit at the index position of v. The following schedule describes which nodes flip the index bit at which stage. A node v that does not have a target does not flip this bit at all. A node v whose label is larger than that of its target flips $I(v)$ in stage 1 if $I(v) = 0$, and in stage 2 if $I(v) = 1$. A node v smaller than its target flips $I(v)$ in stage 3 if $I(v) = 0$, and in stage 4 if $I(v) = 1$.

Theorem 3.2. *Algorithm Flip-Labels completes local communication in an n-node ring with knowledge radius* 0 *in time* $O(\log n)$*, in the radio model.*

Proof. Fix a node v and let u and w be its neighbours. Let t be the index of v, *i.e.*, the first position at which u and w differ. We first prove that after Part 1 every node learns the labels of both of its neighbours and learns the message of (at least) one of them. Suppose, without loss of generality, that $u(t) = 0$ and $w(t) = 1$. Suppose, again without loss of generality, that $v(t) = 1$. Then in step t of the first phase u acts as a transmitter and both v and w act as receivers. Hence v learns the message of u and its label. Using the label of u and its own label, node v can now deduce the label of w as follows. Fix a position i of the labels. If $u(i) = 0$ and $v(i) = 0$ then $w(i) = 1$ if and only if v received a message from w in step i of phase 2. If $u(i) = 0$ and $v(i) = 1$ then $w(i) = 1$ if and only if v received a message from u in step i of phase 1. If $u(i) = 1$ and $v(i) = 0$ then $w(i) = 0$ if and only if v received a message from u in step i of phase 2. Finally, if $u(i) = 1$ and $v(i) = 1$ then $w(i) = 0$ if and only if v received a message from w in step i of phase 1.

Next consider Part 2 and suppose that node v did not get a message from one of its neighbours in Part 1. As before, suppose without loss of generality that, for index t of v, we have $u(t) = 0$, $w(t) = 1$, and $v(t) = 1$. Hence, as argued above, v gets the message of u in Part 1, and consequently w becomes the target of v. Also suppose that the label of v is larger than that of w. (All other cases – there are 8 cases altogether – are similar.) According to the schedule of Part 2, v flips $I(v)$ from 1 to 0 in stage 2 of Part 2. Since $u(t) = 0$, node u does not flip its tth bit at this stage. We have to argue that neither does w. Suppose w does flip its tth bit in stage 2. Hence the index of w is also t. Since the label of w is smaller than that of v and w flips a bit in stage 2, it means that v cannot be the target of w, and hence the target of w is its other neighbour w'. Since t is the index of w

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

and $v(t) = 1$ (originally), it means that $w'(t) = 0$ (originally). This implies, however, that w got a message from w' in Part 1, and thus contradicts the fact that w' became a target of w.

It follows that none of the neighbours of v flips its tth bit at stage 2 and node v does. Consequently Part 1 is run at this stage with $u(t) = v(t) = 0$ and $w(t) = 1$. This means that in step t of phase 2 node v will learn the missing message of its target w.

The above analysis shows that Algorithm Flip-Labels is correct, *i.e.*, every node learns messages of both of its neighbours. Its running time is $2k$ rounds for Part 1 and $4 \cdot 2k$ rounds for Part 2, which gives a total of $10k \in O(\log n)$. □

Corollary 3.3. *Distributed constant colouring can be done in time $O(\log n)$, in an n-node ring with knowledge radius 0, in the radio model.*

Proof. First use Algorithm Flip-Labels twice, thus acquiring knowledge radius 2. Then, by Proposition 2.7, distributed constant colouring can be done in time $O(\log \log n)$. □

Our next result establishes a matching lower bound on time of local communication and of distributed constant colouring in the radio model, thus proving that the running times of Algorithm Flip-Labels, and of the algorithm for distributed constant colouring derived from it, have an optimal order of magnitude.

Theorem 3.4. *Every local communication algorithm and every distributed constant colouring algorithm working in an n-node ring, for knowledge radius 0, have running time $\Omega(\log n)$, in the radio model.*

Proof. We first show that every local communication algorithm must use time at least $(\log n)/2$, for sufficiently large n. Suppose that there exists an algorithm which completes this task in time $t < (\log n)/2$. Fix a round $i \leqslant t$ and a node with label v. The knowledge of a node in round i consists of the following: its label v, the integer i, and all messages obtained before round i. Thus the algorithm yields a function f specifying the decision made in round i by a node with label v, under the assumption that the node has not received any message before round i. The value $f(i, v)$ of this function can be 0 or 1, corresponding to the decisions: act as a receiver or act as a transmitter.

Let $A_i(\varepsilon) = \{v : f(i, v) = \varepsilon\}$, for $\varepsilon = 0, 1$. An *atom* is the set $\bigcap_{i=1}^{t} A_i(\varepsilon_i)$, for any binary sequence $(\varepsilon_i : i = 1, \ldots, t)$. Since $t < (\log n)/2$, at least one of the 2^t atoms has size at least $\sqrt{n}$. Choose one such large atom and denote it by C. All nodes in C have the following property: for any round $i \leqslant t$, if they have not received any message before round i then they behave identically: they all act as transmitters, or all act as receivers.

Suppose that all nodes from C are positioned as a segment of the ring, in any permutation. Consider the middle node v of this segment (or one of the two middle nodes). Since, for sufficiently large n, it is at distance larger than t from all nodes outside C, it follows by induction on i that it does not get any message in any round $i \leqslant t$ because its neighbours behave identically in all those rounds. This contradicts the fact that the algorithm completes local communication in time t.

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

We finally prove that distributed constant colouring cannot be completed in time t either. Any algorithm performing this task in this time yields two functions: a function f defined as above, and a function $g : \{1,\ldots,n\} \longrightarrow \{1,\ldots,c\}$, where c is the number of colours, such that $g(v)$ is the colour output by node with label v in round t, assuming that this node has not received any message before round t.

Consider the atom C defined above (using function f). There exist two distinct labels $u, w \in C$ for which $g(u) = g(w)$. Position all nodes from C in a segment, with u and w occupying neighbouring central positions. Similarly as before, since u and w are at distance larger than t from all nodes outside C, none of them gets any message in any round $i \leqslant t$. Consequently they both output colour $g(u) = g(w)$, which is a contradiction because u and w are neighbours. This completes the proof of lower bound $\Omega(\log n)$ on time required for both tasks under consideration. □

4. Conclusion

Our work suggests many research avenues whose aim would be to better understand relations between information available to processors in a distributed system and the efficiency of performing distributed tasks. First, it would be interesting to generalize our study to arbitrary networks: such a generalization could be quite challenging, especially for networks of large degree. Second, other important tasks could be considered in this context, such as broadcasting, information exchange, routing, load balancing, constructing a minimum spanning tree, or a distributed consensus. Third, other types of communication models could be discussed. Finally, information about the network may be not merely local, as captured by the notion of knowledge radius; it may involve such global parameters as the number of nodes in the network, its diameter, or its maximum degree.

References

[1] Albers, S. and Henzinger, M. R. (1997) Exploring unknown environments. In *Proc. 29th Symp. on Theory of Computing*, pp. 416–425.

[2] Alon, N., Bar-Noy, A., Linial, N. and Peleg, D. (1992) Single round simulation on radio networks. *J. Algorithms* **13** 188–210.

[3] Attiya, H., Snir, M. and Warmuth, M. (1988) Computing on an anonymous ring. *J. Assoc. Comput. Mach.* **35** 845–875.

[4] Awerbuch, B., Goldberg, A., Luby, M. and Plotkin, S. (1989) Network decomposition and locality in distributed computation. In *Proc. 30th Symp. on Foundations of Computer Science*, pp. 364–369.

[5] Awerbuch, B., Goldreich, O., Peleg, D. and Vainish, R. (1990) A tradeoff between information and communication in broadcast protocols. *J. Assoc. Comput. Mach.* **37** 238–256.

[6] Bar-Yehuda, R., Goldreich, O. and Itai, A. (1992) On the time complexity of broadcast in radio networks: An exponential gap between determinism and randomization. *J. Comput. Systems Sci.* **45** 104–126.

[7] Chlamtac, I. and Kutten, S. (1985) On broadcasting in radio networks: Problem analysis and protocol design. *IEEE Trans. Communications* **33** 1240–1246.

[8] Chlebus, B. S., Gąsieniec, L., Gibbons, A., Pelc, A. and Rytter, W. (2002) Deterministic broadcasting in ad hoc radio networks. *Distributed Computing* **15** 27–38.

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press

[9] Chrobak, M., Gąsieniec, L. and Rytter, W. Fast broadcasting and gossiping in radio networks. In *Proc. 41st Annual Symp. on Foundations of Computer Science* (FOCS'2000), pp. 575–581.
[10] Cole, R. and Vishkin, U. (1986) Deterministic coin tossing and accelerating cascades: Micro and macro techniques for designing parallel algorithms. In *Proc. 18th Symp. on Theory of Computing*, pp. 206–219.
[11] Diks, K., Kranakis, E., Malinowski, A. and Pelc, A. (1995) Anonymous wireless rings. *Theoretical Computer Science* **145** 95–109.
[12] Even, S. and Monien, B. On the number of rounds necessary to disseminate information. In *Proc. 1st ACM Symp. on Parallel Algorithms and Architectures* (June 1989), pp. 318–327.
[13] Feige, U., Peleg, D., Raghavan, P. and Upfal, E. (1990) Randomized broadcast in networks. *Random Struct. Alg.* **1** 447–460.
[14] Fraigniaud, P. and Lazard, E. (1994) Methods and problems of communication in usual networks. *Disc. Appl. Math.* **53** 79–133.
[15] Fujita, S., Perennes, S. and Peters, J. G. (1998) Neighbourhood gossiping in hypercubes. *Parallel Processing Letters* **8** 189–195.
[16] Gallager, R. (1985) A perspective on multiaccess channels. *IEEE Trans. Inform. Theory* **31** 124–142.
[17] Goldberg, A. V., Plotkin, S. A. and Shannon, G. E. (1987) Parallel symmetry-breaking in sparse graphs. In *Proc. 19th Symp. on Theory of Computing*, pp. 315–324.
[18] Hedetniemi, S. M., Hedetniemi, S. T. and Liestman, A. L. (1988) A survey of gossiping and broadcasting in communication networks. *Networks* **18** 319–349.
[19] Krumme, D. W. (1992) Fast gossiping for the hypercube. *SIAM J. Comput.* **21** 365–380.
[20] Kutten, S. and Peleg, D. (1995) Fault-local distributed mending. In *Proc. 14th ACM Symp. on Principles of Distributed Computing*, pp. 20–27.
[21] Linial, N. (1992) Locality in distributed graph algorithms. *SIAM J. Comput.* **21** 193–201.
[22] Luby, M. (1988) Removing randomness in parallel computation without a processor penalty. *Proc. 29th IEEE Symp. on Foundations of Computer Science*, pp. 162–173.
[23] Mayer, A., Naor, M. and Stockmeyer, L. (1995) Local computations on static and dynamic graphs. In *Proc. 3rd. Israel Symp. on Theory and Computing Systems.*
[24] Panconesi, A. and Srinivasan, A. (1996) On the complexity of distributed network decomposition. *J. Algorithms* **20** 356–374.
[25] Panaite, P. and Pelc, A. (1999) Exploring unknown undirected graphs. *J. Algorithms* **33** 281–295.

https://doi.org/10.1017/S0963548303006011 Published online by Cambridge University Press