


# Distributed Networks



all our favourite questions:

- can alice send bob a message?
- what is the shortest route?
- colouring?
- (- minimum cut, bipartiteness, etc.)

## Model

Synchronised rounds.

Each round, each vertex independently:

- check inbox, read messages,
- perform unlimited computation,
- message neighbours

Each vertex has unlimited storage, a unique ID (known beforehand) and a list of neighbours (known beforehand)

Congest model:  $O(\log n)$  bit limit



# Colouring a path


Reminder, colouring:

Each vertex is given a colour,  $c: V \rightarrow \{1, \dots, k\}$ , such that each edge  $(v, u)$  is bichromatic, ie.  $c(v) \neq c(u)$

RAM-model colouring of a path: 1-2-1-2-1-2-... trivial! But in the distributed model? Less trivial.

## Colouring algorithm for 3-colouring a path

Assume every vertex has a unique identifier of  $O(\log n)$  bits.





Round 1: send ID to neighb.s.

All other rounds:

- read neighbour's IDs from inbox
- if  $\begin{cases} \text{myID} > \max \{ \text{ID of neighbour} \} \\ \text{myID} > 3 \end{cases}$

then: choose  $ID \in \{1, 2, 3\}$ .


send ID to neighbours





# Breadth-first search tree

Reminder: Given a graph and a source, we may divide into layers of incr. dist. from s.




BFS connects the next layer to the previous.

RAM-model: Starting with  $L_0 = \{s\}$ , find  $L_i$  as  $\bigcup_{v \in L_{i-1}} \text{neighb. } v$ . Link  $u \in L_i$  to an arbitrary neighbour in  $L_{i-1}$ .

## Wave algorithm




Given a source, s, who initiates the wave (everyone else passive)



Round 1: s sends message to all its neighb.s switches state to ~~active~~ "done".

All other rounds: for each  $v \in \text{Network}$ : if v is not done:

- check messages.
- if there is at least 1 message let u be the neighbour who sent that message (choose arbit. vary.) set  $p(v) = u$ . mark self as "done".
- send message to all neighb.s.




Result: every vertex con. to s can send s a message via a shortest path:  $v \rightarrow p(v) \rightarrow p(p(v)) \rightarrow \dots \rightarrow s$ .



Wave with distances: small modification.



Round 1: s sends 0 to nbs.



Round > 1: let  $i$  be the message content received in the arbitrarily chosen msg. send  $i+1$  to neighbors.

Result: every vertex knows its distance to  $s$ .



Problem: Who are my children? When are we <sup>all</sup> done?

Breadth-first search tree algorithm

$v \in \text{Network}$  has extra information:

$d(v)$  distance to root.

$p(v)$  parent in BFS-tree

$C(v) = \{c_1, c_2, \dots\}$  children in BFS-tree

$a(v)$  are we done? - bit. Namely:

are all my children's subtrees done.

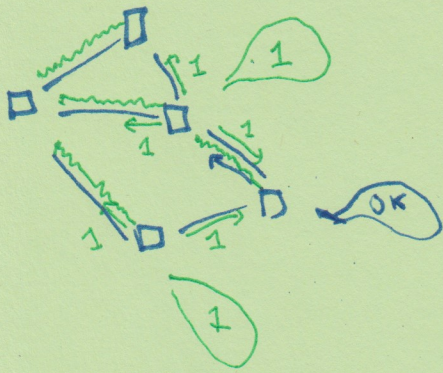


$p(v) = \dots$   
 $d(v) = i$   
 $C(v) = \{c_1, c_2\}$   
 $a(v) = \text{false}$   
 (or true)



# (BFS-tree algorithm, contd)

(5)



Modification 1: Accepting a parent.

When setting  $p(v) = u$ , notify  $u$ .

Result:  $u$  knows  $v \in C(u)$ .

Modification 2: Are we done yet - bits.


if  $C(v)$  has been determined

my  $a(v)$  is the logical and

$$a(c_1) \wedge a(c_2) \wedge \dots \text{ for } c_i \in C(v)$$

thus: keep booleans  $\{a_1, a_2, \dots\}$


- read messages containing "ack" of children,
- update booleans.
- if now all have turned "true" send "ack" to  $p(v)$ .



Result:  $s$  knows ~~when~~ at some point that her message has reached every vertex.


Round complexity: Let  $z$  be a vertex

of maximal distance to  $s$ . It takes  $d(s, z)$  rounds to reach  $z$ . Then  $O(1)$  rounds for  $z$  to realize  $C(z) = \emptyset$ . Then  $d(s, z)$  rounds to send back an "ack". Total:  $O(\text{diam}(G))$





BFS-tree recap: the life of  $v \in L_i$ . (6)



for rounds  $1, 2, \dots, i-1$ : nothing happens

round  $i$ :  $d(v) := i$ ,  $p(v) := u$ ,  $u \in L_{i-1}$

round  $i+1$ :  $C(v) := \{c_1, c_2, \dots\}$

rounds  $i+1$  to  $x$ : waiting to hear back from children.

round  $x+1$ : send ack to  $p(v)$ .

further rounds: nothing happens.

Round analysis: Notice that  $x$  can be no larger than  $d+1+(d-i)$ .

Leader election: finding the minimum ID in  $G$ .

Round 1: everyone starts BFS from themselves.

Round  $> 1$ : • check msg and update min. ID

• proceed as BFS from source min ID

• augment messages with min ID.

Result: • for  $s$  with min. ID in  $G$ , BFS terminates with ack from all children.

• for  $v \neq s$ , there is  $c_i \in C(v)$  who never sends "ack". Thus, every vertex knows whether they lead.







# Token walk analysis (n is the num. of vertices) (8)

Round complexity: ~~every~~ every tree-edge is traversed

$O(n)$  rounds.

$\Leftarrow$

2 times. there are  $n-1$  tree-edges.  
there is one visit per move:

$O(n)$  token movements.

After the last token movement,  
it takes  $O(\text{diam}(G))$  rounds to finish.

## Correctness:

We boldly claimed that vertex  $v$  received  
at most 1 distance-update in each round.

Let us prove this statement by contradiction:

assume  $v$  receives the wave from both  $x$  and  $y$   
for the first time in round  $i$ .

Assume  $\text{Wave}(x)$  was initiated before  $\text{Wave}(y)$ ,  
at times  $t_x$  and  $t_y$ , respectively.

Then, our distance to  $x$  is  $i - t_x$  and to  $y$  is  $i - t_y$ .

But because the token is slow,  $t_y - t_x = 2 \cdot \text{dist}(x, y)$ .

$$i - t_x = \text{dist}(x, v) \leq \text{dist}(y, x) + \text{dist}(y, v) = \text{dist}(y, x) + i - t_y$$

$$\Rightarrow t_y - t_x \leq \text{dist}(x, y). \quad \Rightarrow 2 \cdot \text{dist}(x, y) \leq \text{dist}(x, y)$$

$\Rightarrow \text{dist}(x, y) = 0$  and thus  $x = y$ , a contradiction.  $\square$