### Self-stabilizing Spanning Tree

Peva BLANCHARD

-

- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

æ

< Ē

- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

3/30

#### Remember?



#### Remember?



#### Self-stabilizing

### EXIT



5/30

#### Self-stabilizing





#### Self-stabilizing





- nodes = processes  $p_0, \ldots, p_{n-1}$
- *directed* edge (*p<sub>i</sub>*, *p<sub>j</sub>*) : atomic RW register *r<sub>ij</sub>*
- p<sub>i</sub> writes to r<sub>ij</sub>
- p<sub>i</sub> reads from r<sub>mi</sub>



A D > A B >

#### Assumptions

• *p*<sub>0</sub> is a distinguished processor (root) and knows it



э

Assumptions

- p<sub>0</sub> is a distinguished processor (root) and knows it
- each *p<sub>i</sub>* knows an ordered list *N<sub>i</sub>* of its neighbours



#### Register r<sub>ij</sub>

• boolean field parent :

$$p_i$$
 points to  $p_j \iff r_{ij}.parent = 1$ 

#### • integer field *dist* : distance from the root $p_0$ to $p_i$ .



• 1. The problem

#### • 2. The algorithm

- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

9/30

# The algorithm EXIT



# The algorithm EXIT













◆□ ▶ ◆□ ▶ ◆ □ ▶ ◆ □ ▶ ◆ □ ▶

#### See blackboard (or exercice sheet)

・ロト ・四ト ・ヨト ・ヨト

Root p<sub>0</sub>



æ

< □ > < @ > < 注

Non-root p<sub>i</sub>



æ

∢ ≣⇒

Image: A matrix and a matrix

Non-root p<sub>i</sub>



● ▶ ● ●

Image: A matrix and a matrix

Non-root p<sub>i</sub>



æ

3

A B A B A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A
 A

Non-root p<sub>i</sub>



æ

∢ 臣 ▶

.

▲□▶ ▲圖▶ ▲ 厘▶

< ≣⇒

Non-root p<sub>i</sub>



æ

- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

Step at process *p<sub>i</sub>* 

- sequence of instructions until the next access to some register (read or write)
- considered atomic here



#### Execution

- (infinite) interleaving of steps of processes
- · fairness assumption : every process takes infinitely many steps



#### Asynchronous round

 round 1 : smallest prefix in which each process takes at least one step. and so on



Asynchronous round

- round 1 : smallest prefix in which each process takes at least one step.
- round 2 : next segment in which each process takes at least one step and so on



#### Asynchronous round

- round 1 : smallest prefix in which each process takes at least one step.
- round 2 : next segment in which each process takes at least one step
- and so on



- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

### **Proof - Stability**

Configuration

- for each p<sub>i</sub>, values of local variables Ir<sub>mi</sub>, F, dist
- register values r<sub>ij</sub> for all p<sub>i</sub>, p<sub>j</sub>
- program counter : each process may not start at the beginning of the pseudo-code !

### **Proof - Stability**

Legitimate configuration

- encodes a spanning tree rooted at p<sub>0</sub>
- distance values in each r<sub>ij</sub> is correct
- the parent  $p_i$  of  $p_i$  is the first process in  $N_i$  with minimal distance



### **Proof - Stability**

- no process is inclined to change parent or distance
- the set of legitimate configurations is stable

- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

・ロト・日本・ ・ ヨト ・ ヨト

Strategy : from an arbitrary initial configuration

- Each register *r<sub>ij</sub>* eventually holds the correct distance from the root to *p<sub>i</sub>*
- Then, each process *p<sub>i</sub>* selects the first valid parent in *N<sub>i</sub>* ⇒ this yields a legitimate configuration

Strategy : from an arbitrary initial configuration

- Each register *r<sub>ij</sub>* eventually holds the correct distance from the root to *p<sub>i</sub>*
- Then, each process *p<sub>i</sub>* selects the first valid parent in *N<sub>i</sub>* ⇒ this yields a legitimate configuration

#### Lemma

Let  $\Delta$  = max. nb of neighbours. In every 2 $\Delta$  successive rounds, each  $p_i \neq p_0$  performs "one loop in the pseudo-code".



#### Definition

A *floating distance* in configuration *C* is a value in some field  $r_{ij}$ .*dis* that is smaller than the distance from the root to  $p_i$ .



#### Lemma

Let  $E_k$  be the suffix of execution after the first  $\Delta + 4k\Delta$  rounds.

- (Small(k)) For any C ∈ E<sub>k</sub>, if C has a floating distance, then the smallest floating distance in C is ≥ k
- (Dist(k)) For any C ∈ E<sub>k</sub>, the distance values in registers of processes within distance k from the root are correct

By induction :  $E_1$  after  $\Delta + 4\Delta$  rounds (k = 1)

- each distance field is  $\geq$  0 in the first (arbitrary) configuration
- Proc  $p_i \neq p_0$ 
  - during the first  $2\Delta$  rounds, each non-root  $p_i$  computes dist : dist  $\geq 1$
  - in the next  $2\Delta$  rounds, each non-root  $p_i$  writes *dist* to its registers  $r_{ij}$

 $\Rightarrow$  afterwards, always  $r_{ij} \ge 1$  for all *j*. In particular, Small(1) holds

By induction :  $E_1$  after  $\Delta + 4\Delta$  rounds (k = 1)

- each distance field is  $\geq$  0 in the first (arbitrary) configuration
- root *p*<sub>0</sub>
  - first Δ rounds : p<sub>0</sub> writes 0 in every r<sub>0j</sub>
  - in the next 2Δ rounds, each root's neighbour p<sub>j</sub> reads 0 in r<sub>0j</sub>
  - in the next 2∆ rounds, each root's neighbour p<sub>j</sub> writes 1 to their registers distance fields.
  - $\Rightarrow$  Dist(1) holds

Assume Small(k) and Dist(k) hold in  $E_k$ .

- Let *C<sub>k</sub>* first config of *E<sub>k</sub>*
- $m \ge k$  smallest floating distance in  $C_k$
- In the next 4∆ rounds after C<sub>k</sub>, each proc that chooses m as the smallest value assigns m+1 to its distance
- Thus, afterwards, smallest floating distance  $\geq m + 1 \geq k + 1$
- $\Rightarrow$  Small(k+1) holds in  $E_{k+1}$

Assume Small(k) and Dist(k) hold in  $E_k$ .

- Dist(k) holds in *E<sub>k</sub>* : every proc within distance *k* from the root have registers with correct distance
- Let  $p_i$  be proc at distance k + 1
- In the next 4 $\Delta$  rounds after  $C_k$ :  $p_i$  necessarily chooses value k, and assigns k + 1 to its register distance fields.
- $\Rightarrow$  Dist(k+1) holds in  $E_{k+1}$

- 1. The problem
- 2. The algorithm
- 3. Proof
  - a. Stability
  - b. Convergence
- 4. Conclusion

・ロト・日本・ ・ ヨト ・ ヨト

# Conclusion

#### Assumptions

- *n* processes, with a root *p*<sub>0</sub>
- bidirectional communication with rw register r<sub>ij</sub>
- each *p<sub>i</sub>* knows a an ordered list of its neighbours
- each process takes infinitely many steps
- $\Rightarrow$  the algorithm is self-stabilizing