

The Model & Basic Computations

Chapter 1 and 2

The Model

Broadcast

Spanning Tree Construction

Traversal

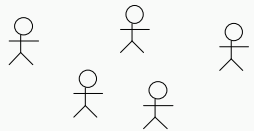
Wake-up



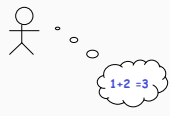
Paola Flocchini

Distributed Environment

Multiplicity



Autonomy

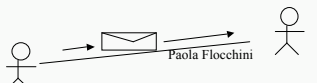


clock

memory

computing capabilities

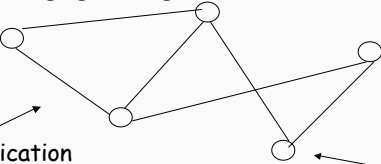
Interaction



typically by exchange of messages

The Model

Collection of entities that communicate by exchanging messages



communication link

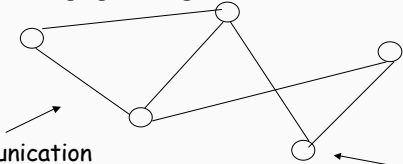
entity, node, site ...

(processor, process, object ...)

Paola Flocchini

The Model

Collection of entities that communicate by exchanging messages



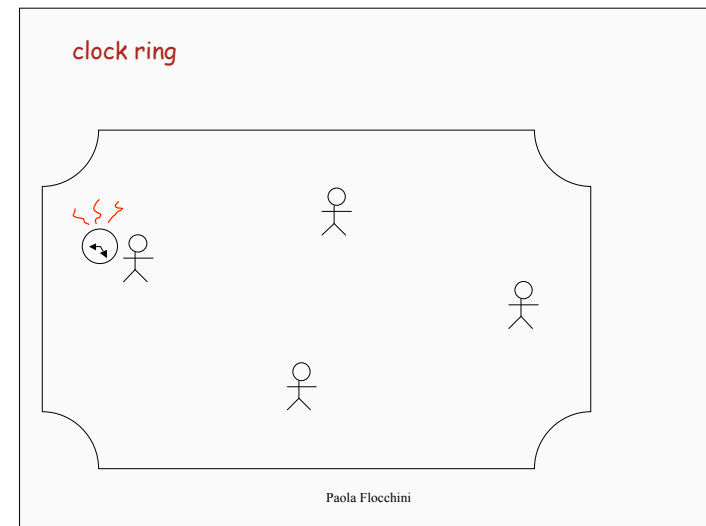
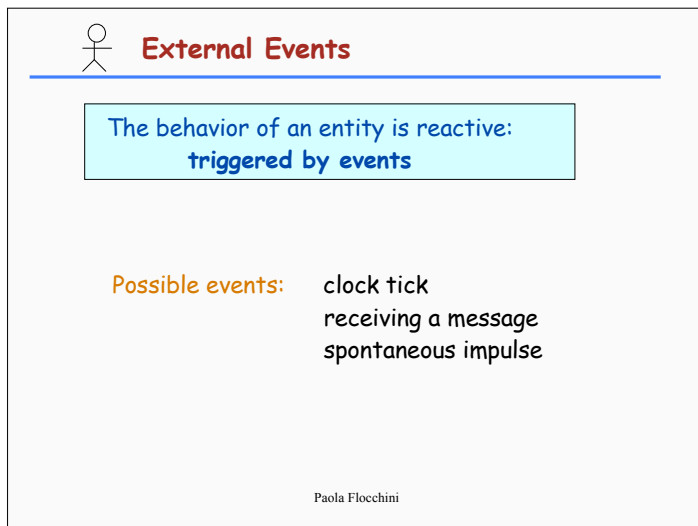
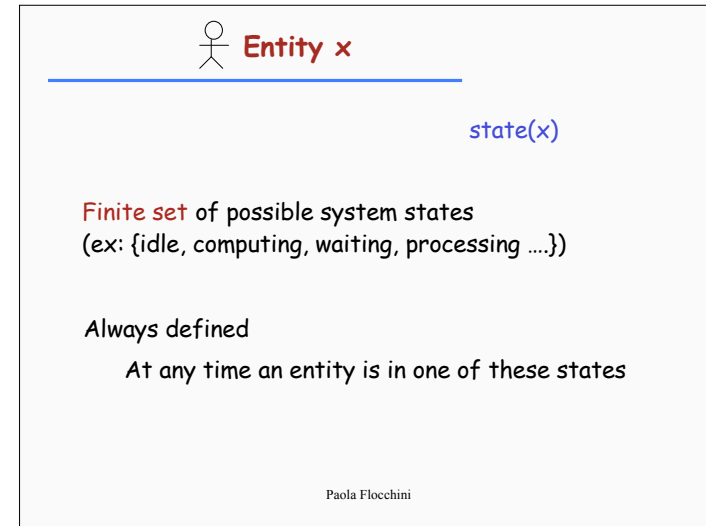
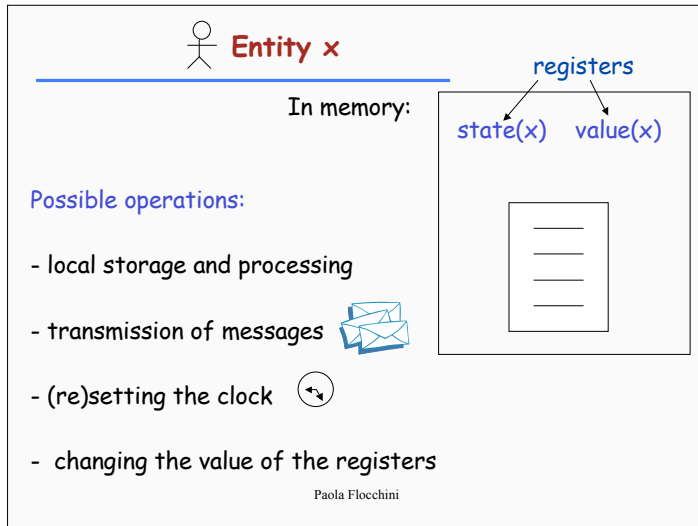
communication link

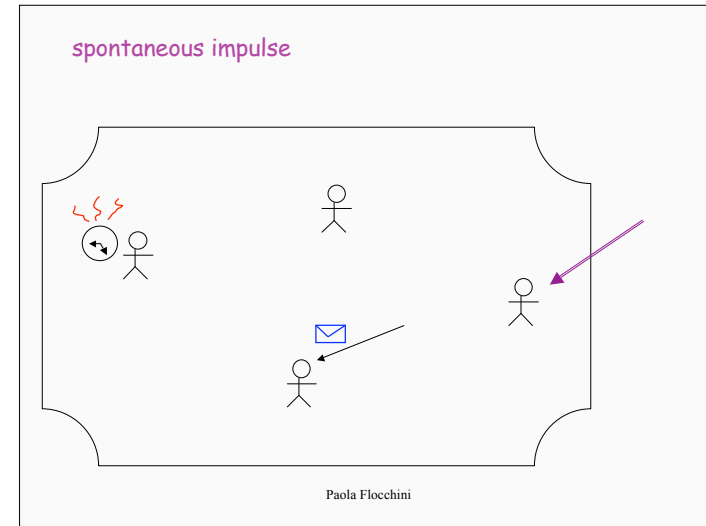
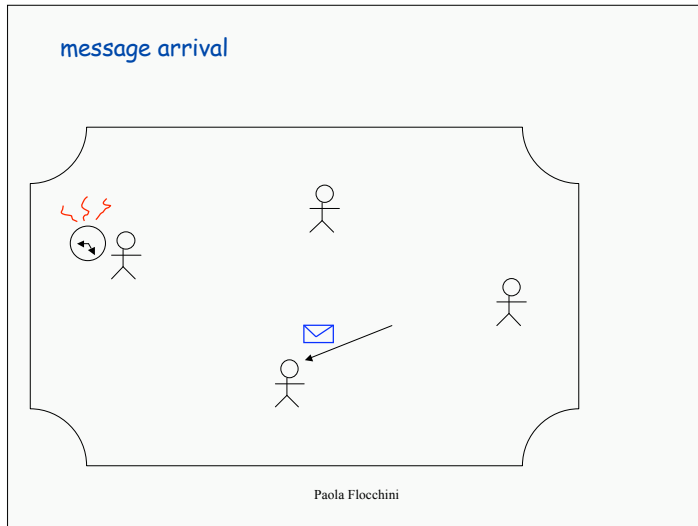
entity, node, site ...

message-passing

(processor, process, object ...)

Paola Flocchini






The reaction of an entity depends on the event and on its state

State x Event \longrightarrow Action

Paola Flocchini

 **Actions**

Action: sequence of activities, e.g.,

- computing
- sending message
- change state

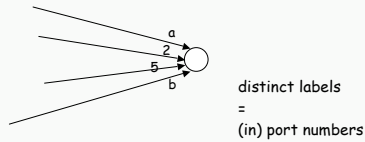
an action is **atomic**
the activities cannot be interrupted

an action is **terminating**
the activities must terminate within finite time

Paola Flocchini

Local orientation: more precisely

Each entity distinguishes among its in-neighbors

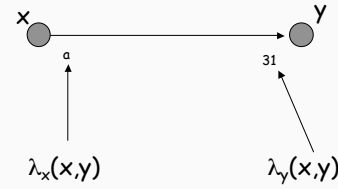


When a message arrives, the entity can detect from which port

Paola Flocchini

Local orientation: more precisely

for an edge (x,y) there are two labels:

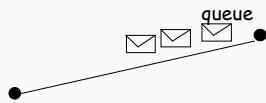


Topology = labeled graph (G, λ)

Paola Flocchini

Restrictions of the model: examples

Communication Restrictions



Message Ordering

(FIFO)

In absence of failures, msgs transmitted along the same link arrive in the same order.

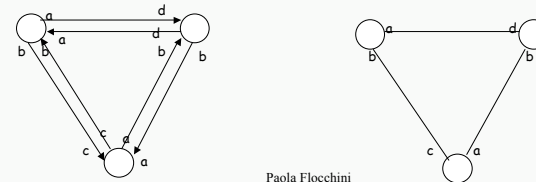
Paola Flocchini

Restrictions of the model: examples

Communication Restrictions

Bidirectional Links

$\forall x, N_i(x) = N_o(x) = N(x)$ and
 $\forall y \in N(x): \lambda_x(x,y) = \lambda_x(y,x)$



Restrictions of the model: examples

Reliability Restrictions:

1. *Guaranteed delivery:*
Any message that is sent will be received uncorrupted
2. *Partial Reliability:*
There will be no failures
3. *Total Reliability:*
No failures have occurred nor will occur

..... Paola Flocchini

Restrictions of the model: examples

Topological restriction:

The graph G is strongly connected

.....

Knowledge Restrictions

Knowledge of number of nodes
Knowledge of number of links
Knowledge of diameter

..... Paola Flocchini

Restrictions of the model: examples

Time restriction:

Bounded Communication Delay:

There exists a constant Δ such that, in absence of failures, the communication delay of any message on any link is at most Δ


Synchronized clocks:

All local clocks are incremented by one unit simultaneously and interval are constant

..... Paola Flocchini

Complexity measures - Performance

1. Amount of communication
number of messages exchanged
(finer granularity: number of bits)



point of view of SYSTEM

2. Time

point of view of USER

Communication delays are in general unpredictable !!!

Ideal time:
1 unit of time to transmit 1 message

Example - Broadcast

Assumptions = Restrictions

- Unique Initiator → By definition of problem
- Total reliability → Simplifying assumptions
- Bidirectional links → Simplifying assumptions
- G is connected → Otherwise unsolvable

Paola Flocchini

Algorithm FLOOD

The idea: If an entity knows something, it sends the info to its neighbours

One entity is INITIATOR, the others are SLEEPING

INITIATOR

spontaneously

send(I) to N(x)

SLEEPING

receiving(I)

send(I) to N(x)

Paola Flocchini

The idea: If an entity knows something, it sends it to its neighbours except the sender

INITIATOR

spontaneously

send(I) to N(x)

SLEEPING

receiving(I)

send(I) to N(x) - {sender}

Paola Flocchini **It does not terminate**

$S = \{\text{initiator, sleeping, done}\}$

Algorithm for node x:

INITIATOR

spontaneously

send(I) to N(x)

become (DONE)

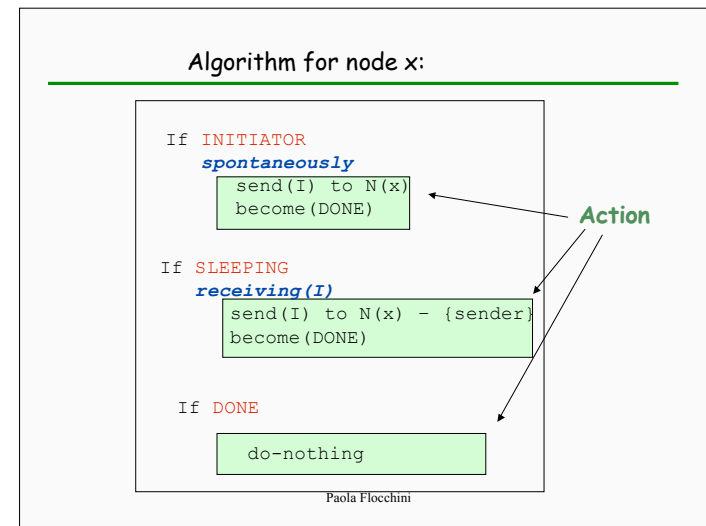
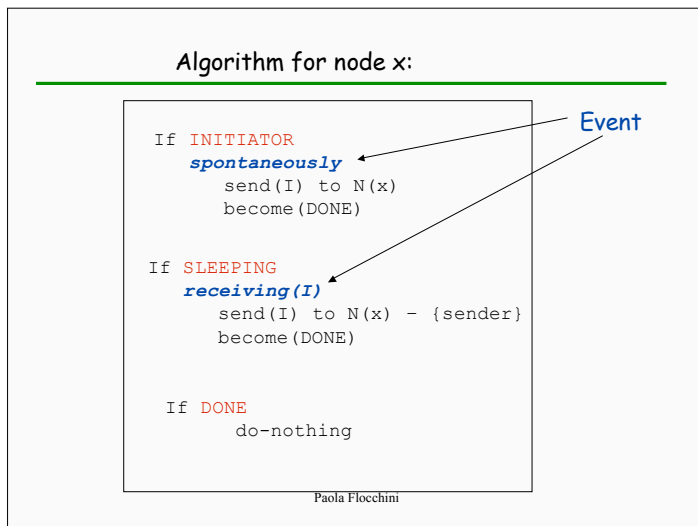
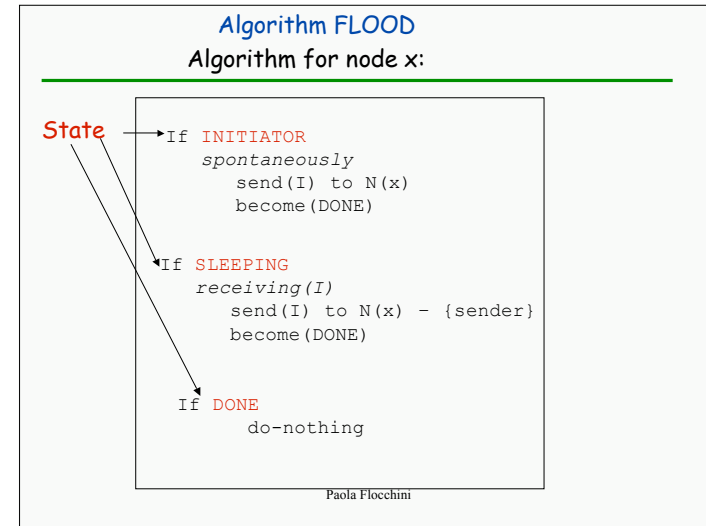
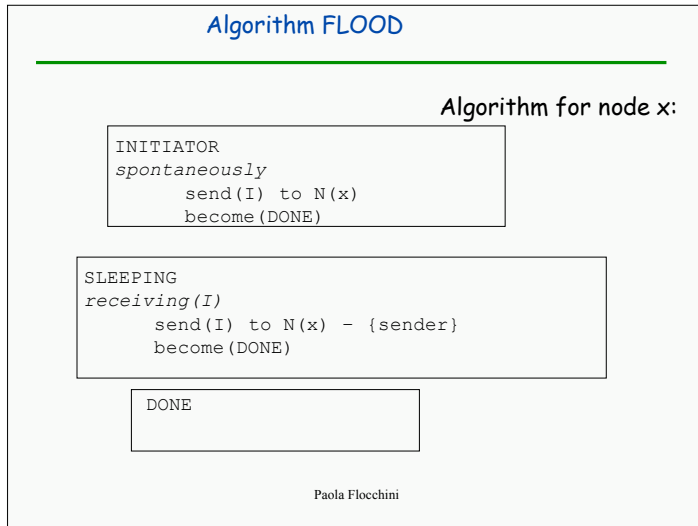
SLEEPING

receiving(I)

send(I) to N(x) - {sender}

become (DONE)

Paola Flocchini



Example

Paola Flocchini

Correctness

The Algorithm terminates in finite time

It follows from: G connected and total reliability

Termination

Local Termination: when DONE

Global Termination: when?

Paola Flocchini

Message complexity Worst Case

m = number of links Worst for all possible initiators
and for all possible executions

Messages: ≤ 2 on each link

$\leq 2m$

$O(m)$

More precisely:

Let s be the initiator

$$|N(s)| + \sum_{x \neq s} (|N(x)| - 1)$$

$$= \sum_x |N(x)| - \sum_{x \neq s} 1$$

$\sum_x |N(x)| = 2m$

$2m - (n-1)$

Paola Flocchini

Time Complexity - Ideal Time Worst Case

Worst for all possible initiators
and for all possible executions

Time: (ideal time)

$$\text{Max}_x \{d(x,s)\} = \text{eccentricity of } s$$

$$\leq \text{Diameter}(G) \leq n-1$$

$O(n)$

Paola Flocchini

Time and Events

External events:
spontaneously
receiving
when (clock)

Actions may generate events

send generates *receiving*
set-clock generates *when*

Generated events might not occur (in case of faults).
 If they occur, they occur *later*.

In the case of *receiving* with some unpredictable delay.

Paola Flocchini

Different delays ----> *different executions*

Different executions could have *different outcomes*

(Spontaneous events are considered generated *before* execution starts: initial events)

An execution is fully described by the *sequence of events* that have occurred

Paola Flocchini

Time x Event diagram

Spontaneous event

receiving event

x

y

time

as seen by an observer

Paola Flocchini

Example: Time x Event diagram of Flooding

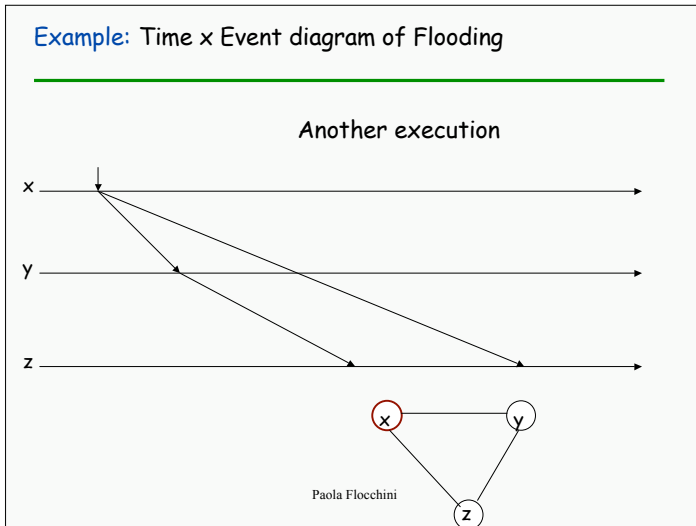
One possible execution

x

y

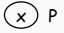
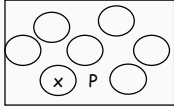
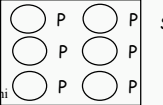
z

Paola Flocchini



Knowledge

$P = \text{fact}; x = \text{entity}; S = \text{set of entities.}$

- **Local knowledge LK**
 $P \in LK(x).$

- **Implicit knowledge IK**
 $P \in IK(S) \text{ if } \exists x \in S: P \in LK(x).$

- **Explicit knowledge EK**
 $P \in EK(S) \text{ if } \forall x \in S: P \in LK(x).$


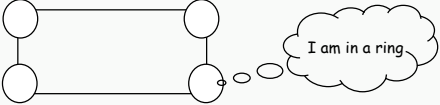
Paola Flocchini

- **Common knowledge CK**
 $P \in CK(S) \text{ if}$
 $\forall x \in S, P \in LK(x) \wedge$
 $\forall x \in S (\forall x \in S, P \in LK(x)) \in LK(x) \wedge$
 $\forall x \in S ((\forall x \in S, P \in LK(x)) \in LK(x)) \in LK(x) \wedge \dots$

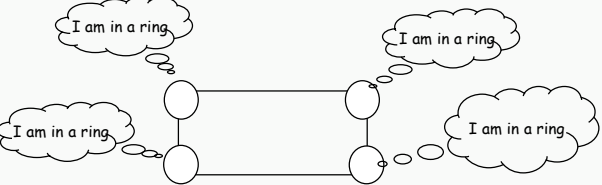
Paola Flocchini

Examples

Implicit knowledge



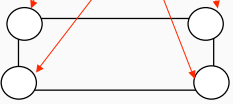
Explicit knowledge



Paola Flocchini

Common knowledge

I know I'm in a ring,
everyone knows it,
everyone knows that
everyone knows it, ...
....

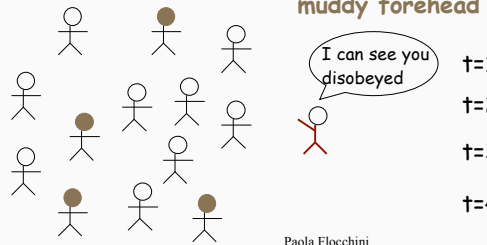


Paola Flocchini

How to reach common knowledge in FINITE TIME ?

$P \in CK(S)$ if
 $\forall x \in S, P \in LK(x) \wedge$
 $\forall x \in S (\forall x \in S, P \in LK(x)) \in LK(x) \wedge$
 $\forall x \in S ((\forall x \in S, P \in LK(x)) \in LK(x)) \in LK(x) \wedge \dots$

muddy forehead



t=1
t=2
t=3
t=4

Paola Flocchini

Some types of knowledge

Topological knowledge
 Graph type ("G is a ring"...), adjacency matrix of G ...

Metric knowledge
 Number of nodes, diameter, eccentricity...

Sense of direction
 Information on link labels
 Information on node labels

As the available knowledge grows, the algorithm becomes less portable (rigid). Generic algorithms do not use any knowledge.

Paola Flocchini

Example: impact of knowledge

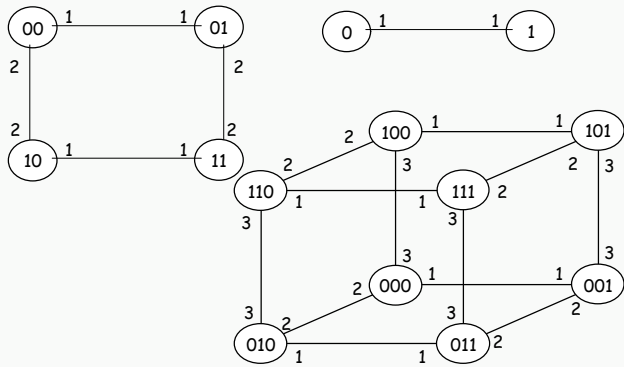
In specific topologies flooding can be avoided and broadcast can be much more efficient (if the topology is known).

What is the complexity of flooding in a complete graph ?
 How can it be done more efficiently ?

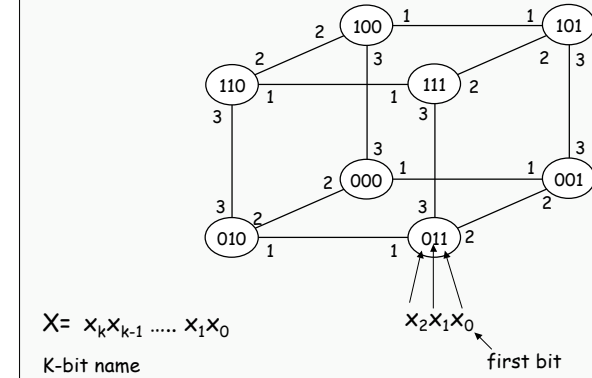
What is the complexity of flooding in a tree ?
 Can it be done more efficiently ?

Paola Flocchini

Example: The labeled hypercube



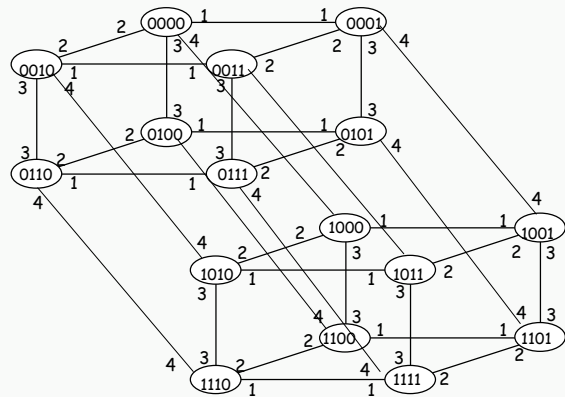
Each link between two nodes is labeled by the dimension of the bit by which the nodes' name differ.



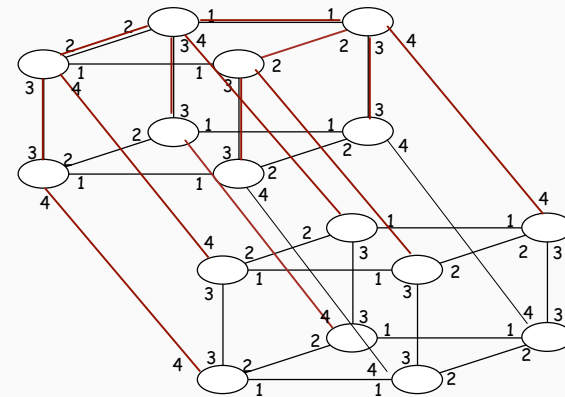
$$X = x_k x_{k-1} \dots x_1 x_0$$

K-bit name

Paola Flocchini



Paola Flocchini



Paola Flocchini

A hypercube of dimension k has $n = 2^k$ nodes

Each node has k links

→ $m = n k/2 = O(n \log n)$

Flooding would cost $O(n \log n)$

Paola Flocchini

HyperFlood - Efficient Broadcast

- 1) The initiator sends the message to all its neighbours
- 2) A node receiving the message from link l , sends it only to links with label $l' < l$

Paola Flocchini

Correctness

Every node is touched

Based on the lemma:

For each pair of nodes x and y there exists a path of decreasing labels

$X = x_k, x_{k-1}, \dots, x_1, x_0$

$Y = y_k, y_{k-1}, \dots, y_1, y_0$

Paola Flocchini

Correctness

Every node is touched

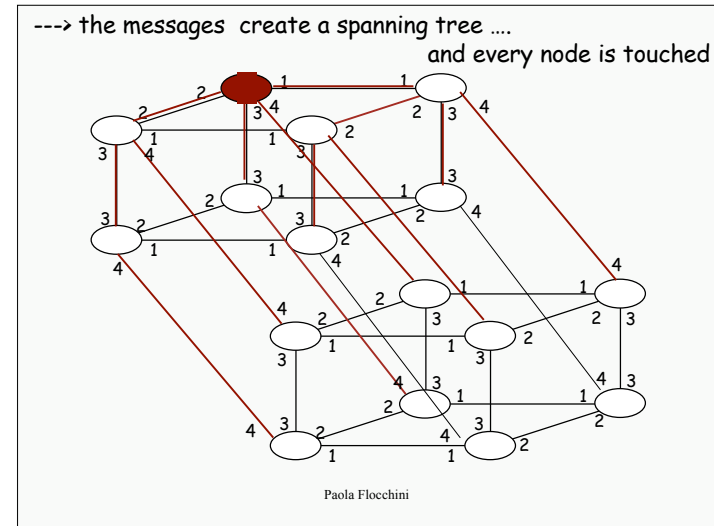
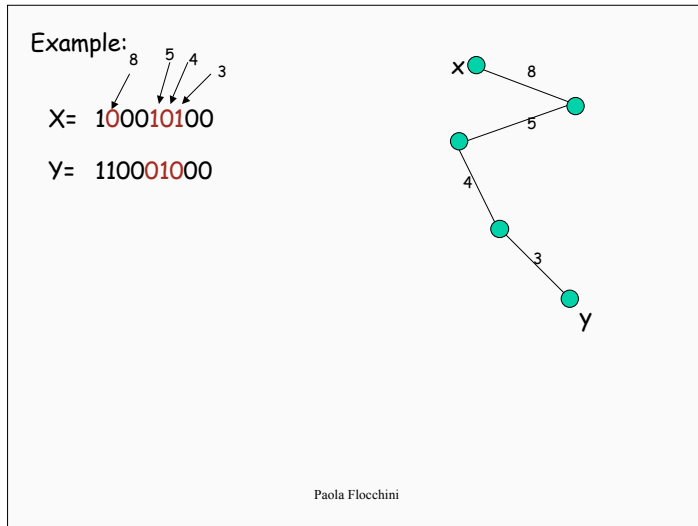
Based on the lemma:

For each pair of nodes x and y there exists a path of decreasing labels

$X = x_k, x_{k-1}, \dots, x_1, x_0$ $Y = y_k, y_{k-1}, \dots, y_1, y_0$

Consider positions where they differ in decreasing order ...

Paola Flocchini



Complexity: $n-1$ (OPTIMAL)

Because every entity receives the info only ONCE.

Paola Flocchini

In Special Topologies

General Flooding: $2m - (n-1)$

Ad-hoc algorithm in hypercube: $(n-1)$

Ad-hoc algorithm in complete network: $(n-1)$

In the tree Flooding is optimal: $(n-1)$

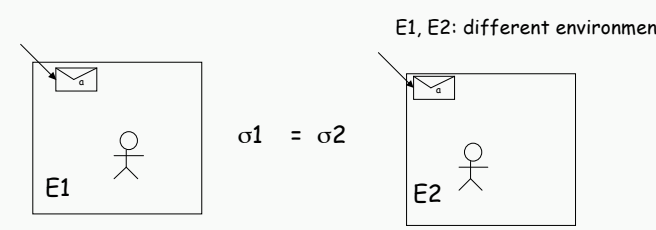
Paola Flocchini

Important Facts

State x Event ----> Action

$\sigma(x,t)$ = internal state of entity x
content of memory (registers, clock, ...) at time t

Paola Flocchini

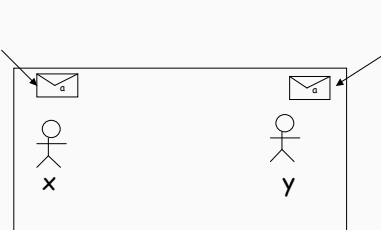


E1, E2: different environments

$\sigma_1 = \sigma_2$

1) If the same event happens to x at time t in two different executions and if the internal states σ_1 and σ_2 of x in the two executions at that time are equal, then **the new internal state of x will be the same in both executions**

Paola Flocchini



$\sigma(x) = \sigma(y)$

2) If the same event happens to x and y at time t in the same execution and if the internal states $\sigma(x)$ and $\sigma(y)$ are equal, then **the new internal states of x and y will be the same.**

Paola Flocchini

An example

Back to **Broadcast** ...

Theorem: Under the set of assumptions:

- unique Initiator
- G is connected
- no failures
- bidirectional links

Every generic broadcast protocol requires, in the worst case, m messages.

Paola Flocchini

Lower Bounds for Broadcast

Proof.

$$m(G) = n. \text{ of edges in } G$$

By contradiction.

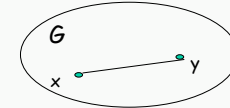
Let A be an algorithm that broadcasts exchanging **less than** $m(G)$ messages (in all executions, and for any graph G) under those assumption.

Then there is at least a link in G where no messages are sent.

Paola Flocchini

Let $e = (x,y)$ be such a link.

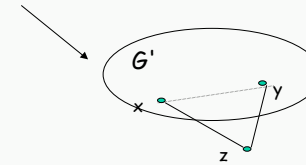
$$G=(V,E)$$



Construct a new graph G'

(remember: n is unknown)

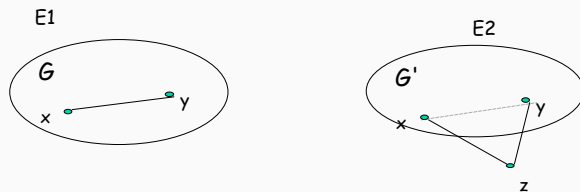
$$G'=(V \cup \{z\}, E-e \cup \{(x,z),(y,z)\})$$



Execute the same algorithm on G' with the same time delays, same initial internal states for all nodes except for z which is sleeping

Paola Flocchini

Two executions in two environments



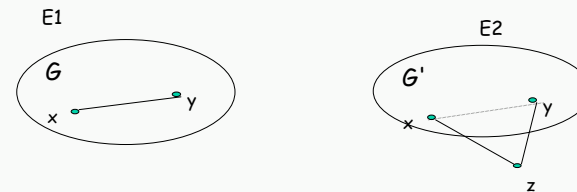
For all nodes, except z , the two executions are **identical**

x and y never send to each other in E1

--->

x and y never send to z in E2

Paola Flocchini



Within finite time the protocol terminates

but in E2 node z will never be reached.

Paola Flocchini

Observations:

1) Dense networks = more messages
(ex. in complete networks $m = n(n-1) \dots$)

2) It is optimum in acyclic graphs

Idea: to solve broadcast.

1. Build a spanning tree of G

2. Execute flooding



Spanning Tree construction Problem

Paola Flocchini