

# Two stories about love for Luis Montejano's 60th birthday

(A topological perspective on distributed computing)

Sergio Rajsbaum  
Instituto de Matemáticas  
UNAM

*From the book coauthored with Maurice Herlihy  
and Dmitry Kozlov to be published by Elsevier*

# The setting

- Characters: men and woman
- Communication maybe limited or unreliable
- Each one has partial information about reality

*Yet, the characters need to solve some task*

# Concurrency is confusing

It may be easy to follow *sequential* procedures, such as preparing an omelette from a recipe

It is much harder to pursue *concurrent* activities, such as preparing a ten-course meal with limited pots and pans, all while speaking to a friend on the telephone.

# Concurrency is everywhere

Nearly every activity in our society  
depends on the Internet

# Concurrency is everywhere

At a smaller scale:

- as processor feature sizes shrink, they become harder to cool, manufacturers have given up trying to make processors faster.
- Instead, they have focused on making processors more parallel.

# Concurrency is everywhere

At the other extreme:

- Internet, cloud computing and peer-to-peer systems may encompass thousands of machines that span every continent.

# The stories

- Cheating wives  
(A.k.a. muddy children, from knowledge theory)
- Two insecure lovers  
(A.k.a. Coordinated attack, from databases and networking)

This talk is about

# Theory of concurrency

This talk is about

# Theory of concurrency

Using topology

# Distributed systems...

- Individual sequential processes (men and women)
- Cooperate to solve some task
- By message passing, shared memory, or any other mechanism

# ... and topology

Many models, appear to have little in common besides the common concern with complexity, failures and timing.

*Combinatorial topology provides a common framework that unifies these models.*

# How topology

“freeze” all possible interleavings and failure scenarios into a single, static, simplicial complex

Appeal of the approach:

reduces the problem of reasoning about computations that unfold in time to the more familiar problem of reasoning about static combinatorial structures.

# Why topology

Distributed algorithms are more challenging than their sequential counterparts because each process, has only a limited view of the world (overall state of the computation)

Placing together all these views yields a simplicial complex

**Cheating wives**

# Cheating wives

# Cheating wives

- There were one million married couples.

# Cheating wives

- There were one million married couples.
- 40 wives were unfaithful

# Cheating wives

- There were one million married couples.
- 40 wives were unfaithful
- Each husband knew whether other men's wives were unfaithful but he did not know whether his wife was unfaithful.

# Cheating wives

- There were one million married couples.
- 40 wives were unfaithful
- Each husband knew whether other men's wives were unfaithful but he did not know whether his wife was unfaithful.
- The King of the country announced “There is at least one unfaithful wife” and publicized the following decree

# Cheating wives decree

He asks the following question over and over:

can you tell for sure whether or not you are a cuckold?

# Cheating wives decree

He asks the following question over and over:

can you tell for sure whether or not you are a cuckold?

Assuming that all of the men are intelligent, honest, and answer simultaneously, what will happen?

# Analysis of the puzzle

First  
operational,  
then  
combinatorial

# Operational analysis (I)

First, suppose that exactly one is cuckold

# Operational analysis (I)

First, suppose that exactly one is cuckold

- He sees nobody else, can conclude that he is the one

# Operational analysis (I)

First, suppose that exactly one is cuckold

- He sees nobody else, can conclude that he is the one
- The others cannot tell whether or not they are cuckolds

# Operational analysis (I)

First, suppose that exactly one is cuckold

- He sees nobody else, can conclude that he is the one
- The others cannot tell whether or not they are cuckolds
- At the first question, exactly one says “yes”

# Operational analysis (I)

First, suppose that exactly one is cuckold

- He sees nobody else, can conclude that he is the one
- The others cannot tell whether or not they are cuckolds
- At the first question, exactly one says “yes”
- At the second, all others say “no”

# Operational analysis (2)

Now, suppose that exactly two are cuckolds

# Operational analysis (2)

Now, suppose that exactly two are cuckolds

- They know at least two are cuckolds, because nobody spoke in first round

# Operational analysis (2)

Now, suppose that exactly two are cuckolds

- They know at least two are cuckolds, because nobody spoke in first round
- They see only one cuckold

# Operational analysis (2)

Now, suppose that exactly two are cuckolds

- They know at least two are cuckolds, because nobody spoke in first round
- They see only one cuckold
- At the second question, exactly two says “yes”

# Operational analysis (2)

Now, suppose that exactly two are cuckolds

- They know at least two are cuckolds, because nobody spoke in first round
- They see only one cuckold
- At the second question, exactly two says “yes”
- At the third, all others say “no”

# Operational analysis (3)

Suppose that exactly  $k$  are cuckolds, by induction...

# Operational analysis (3)

Suppose that exactly  $k$  are cuckolds, by induction...

- At the  $k$ -th question, exactly  $k$  say “yes”

# Operational analysis (3)

Suppose that exactly  $k$  are cuckolds, by induction...

- At the  $k$ -th question, exactly  $k$  say “yes”
- At the  $(k+1)$ -th, all others say “no”

# Combinatorial analysis

Local states

# Combinatorial analysis

## Local states

- A local state is a man's state of knowledge

# Combinatorial analysis

## Local states

- A local state is a man's state of knowledge
- It is represented by a vector: in position  $i$  has 0 if man  $i$  is known to be clean, and 1 if cuckold


# Combinatorial analysis

## Local states

- A local state is a man's state of knowledge
- It is represented by a vector: in position  $i$  has 0 if man  $i$  is known to be clean, and 1 if cuckold
- Because man  $i$  does not know its own status, its input vector has  $\perp$  in position  $i$

# Global inputs


Each possible input configuration is represented as a simplex, linking compatible states for the men



meaning that the men can be in these states together

○ 2   ● 1   ○ 3

# Initial Complex




○ 2   ● 1   ○ 3

# Initial Complex

Man 1 knows that man 2 is clean and man 3 is cuckold

no cuckolds




all cuckolds

- - 
  - ◐
- 2      1      3

# Initial Complex

no cuckolds




he may be clean...

all cuckolds

- - 
  -
- 2      1      3

# Initial Complex




- - 
  -
- 2      1      3

# Initial Com

disappears when announced  
“at least one cuckold”

no cuckolds



all cuckolds

that is, men know that each 2-simplex is a possible initial state, except for the one where all are clean

○

2


●

1


◐

3


# Evolution



11:59 AM



12:01 PM



1:01 PM



2:01 PM

2




3

2


1

3


# Evolution



11:59 AM



12:01 PM



1:01 PM



2:01 PM

○

2

●


1

○


3

# Evolution

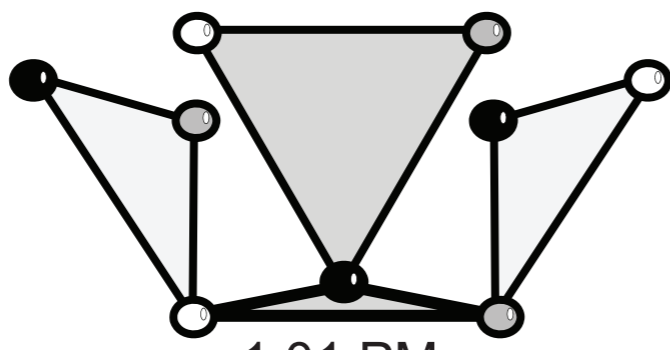
3 vertexes exposed, where someone knows its status



11:59 AM



12:01 PM





1:01 PM




2:01 PM

○ 2   ● 1   ◐ 3

# Evolution



Nobody spoke previous round,  
6 vertexes exposed





2




1




3


# Evolution



11:59 AM



12:01 PM



1:01 PM

All 3 announce "cuckolds"



2:01 PM



2




1




3


# Evolution



11:59 AM



12:01 PM



1:01 PM



2:01 PM

2




3

2


1

3


# Decisions



11:59 AM



12:01 PM



1:01 PM



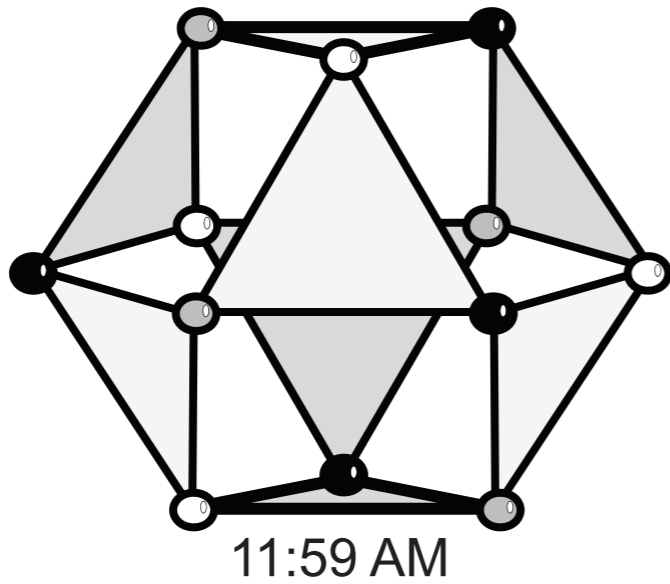


2:01 PM

- 2
- 1
- ◐ 3

# Decisions

No decisions



2



3


2

1


3

# Decisions


3 vertexes  
labeled, "cuckold"



11:59 AM



12:01 PM




1:01 PM




2:01 PM

- 2
- 1
- ◐ 3

# Decisions




11:59 AM



12:01 PM

Nobody spoke previous round,  
6 vertexes labeled "cuckold"



1:01 PM



2:01 PM

2




3

2


1

3


# Decisions



11:59 AM



12:01 PM



1:01 PM



2:01 PM

3 vertexes  
labeled "cuckold"

2




3

2


1

3


# Decisions



11:59 AM



12:01 PM




1:01 PM



2:01 PM

# Output complex


Decisions induce a map to this complex



Each man should say “yes” or “no”  
All combinations are possible...

# Output complex

Decisions induce a map to this complex



... except all “no” after King’s announcement

Each man should say “yes” or “no”  
All combinations are possible...

# Solving the cheating wives task

Each man decides an output value,  
on one of its local states

Decisions define a simplicial map from  
input complex to output complex that  
respects the task's specification

In this task communication is very  
limited. More generally, for any task...

# Solving any task

In the basic, wait-free model

A task is solvable if and only if there exists a *subdivision* of the input complex and a simplicial map to the output complex that respects the task's specification

Herlihy, Shavit 1993

*Wait-free*: asynchronous model where any number of processes can crash

**Two insecure lovers**

# Coordination

We often need to ensure that two things happen together or not at all.

For example, a banking system needs to ensure that if an automatic teller dispenses cash, then the corresponding account balance is debited, and vice-versa.

**Two insecure lovers**

# Two insecure lovers

- Alice and Bob want to schedule a meeting.

# Two insecure lovers

- Alice and Bob want to schedule a meeting.
- If both attend, they win, but if only one attends, defeat and humiliation is felt.

# Two insecure lovers

- Alice and Bob want to schedule a meeting.
- If both attend, they win, but if only one attends, defeat and humiliation is felt.
- As a result, neither will show up without a guarantee that the other will show up at the same time.

# Two insecure lovers

- Alice and Bob want to schedule a meeting.
- If both attend, they win, but if only one attends, defeat and humiliation is felt.
- As a result, neither will show up without a guarantee that the other will show up at the same time.
- Communication is by SMS only.

# Communication problems

- Normally, it takes a message one hour to arrive.
- However, it is possible that it is gets lost.

# The puzzle

Fortunately, on this particular night,  
all the messages arrive safely.

How long will it take Alice and Bob  
to coordinate their meeting?

# Analysis of the puzzle

First  
operational,  
then  
combinatorial

# Operational analysis (I)

Suppose Alice initiates the communication

# Operational analysis

(I)

# Operational analysis

## (I)

- Suppose Bob receives a message at 1:00 from Alice saying “meet at midnight”. Should Bob show up?

# Operational analysis

(I)

- Suppose Bob receives a message at 1:00 from Alice saying “meet at midnight”. Should Bob show up?
- Although her message was in fact delivered, Alice does not know. She therefore considers it possible that Bob did not receive the message.

# Operational analysis

(I)

- Suppose Bob receives a message at 1:00 from Alice saying “meet at midnight”. Should Bob show up?
- Although her message was in fact delivered, Alice does not know. She therefore considers it possible that Bob did not receive the message.
- Hence Alice cannot decide to show up, given her current state of knowledge.

# Operational analysis

(I)

- Suppose Bob receives a message at 1:00 from Alice saying “meet at midnight”. Should Bob show up?
- Although her message was in fact delivered, Alice does not know. She therefore considers it possible that Bob did not receive the message.
- Hence Alice cannot decide to show up, given her current state of knowledge.
- Knowing this, Bob will not show up based solely on Alice’s message.

# Operational analysis

(2)

# Operational analysis

## (2)

- Naturally, Bob reacts by sending an acknowledgment back to Alice, which arrives at 2:00

# Operational analysis

## (2)

- Naturally, Bob reacts by sending an acknowledgment back to Alice, which arrives at 2:00
- Will Alice plan to show up?

# Operational analysis

(2)

- Naturally, Bob reacts by sending an acknowledgment back to Alice, which arrives at 2:00
- Will Alice plan to show up?
- Unfortunately, Alice's predicament is similar to Bob's predicament at 1:00, she cannot yet decide to show up

No number of successfully delivered acknowledgments will be enough to ensure that show up safely!

The key insight is that the difficulty is not caused by what actually happens (all messages actually arrive) but by the uncertainty regarding what might have happened.

# Combinatorial analysis

# Combinatorial analysis

- Initially Alice has two possible decisions: meet at dawn, or meet at noon the next day.

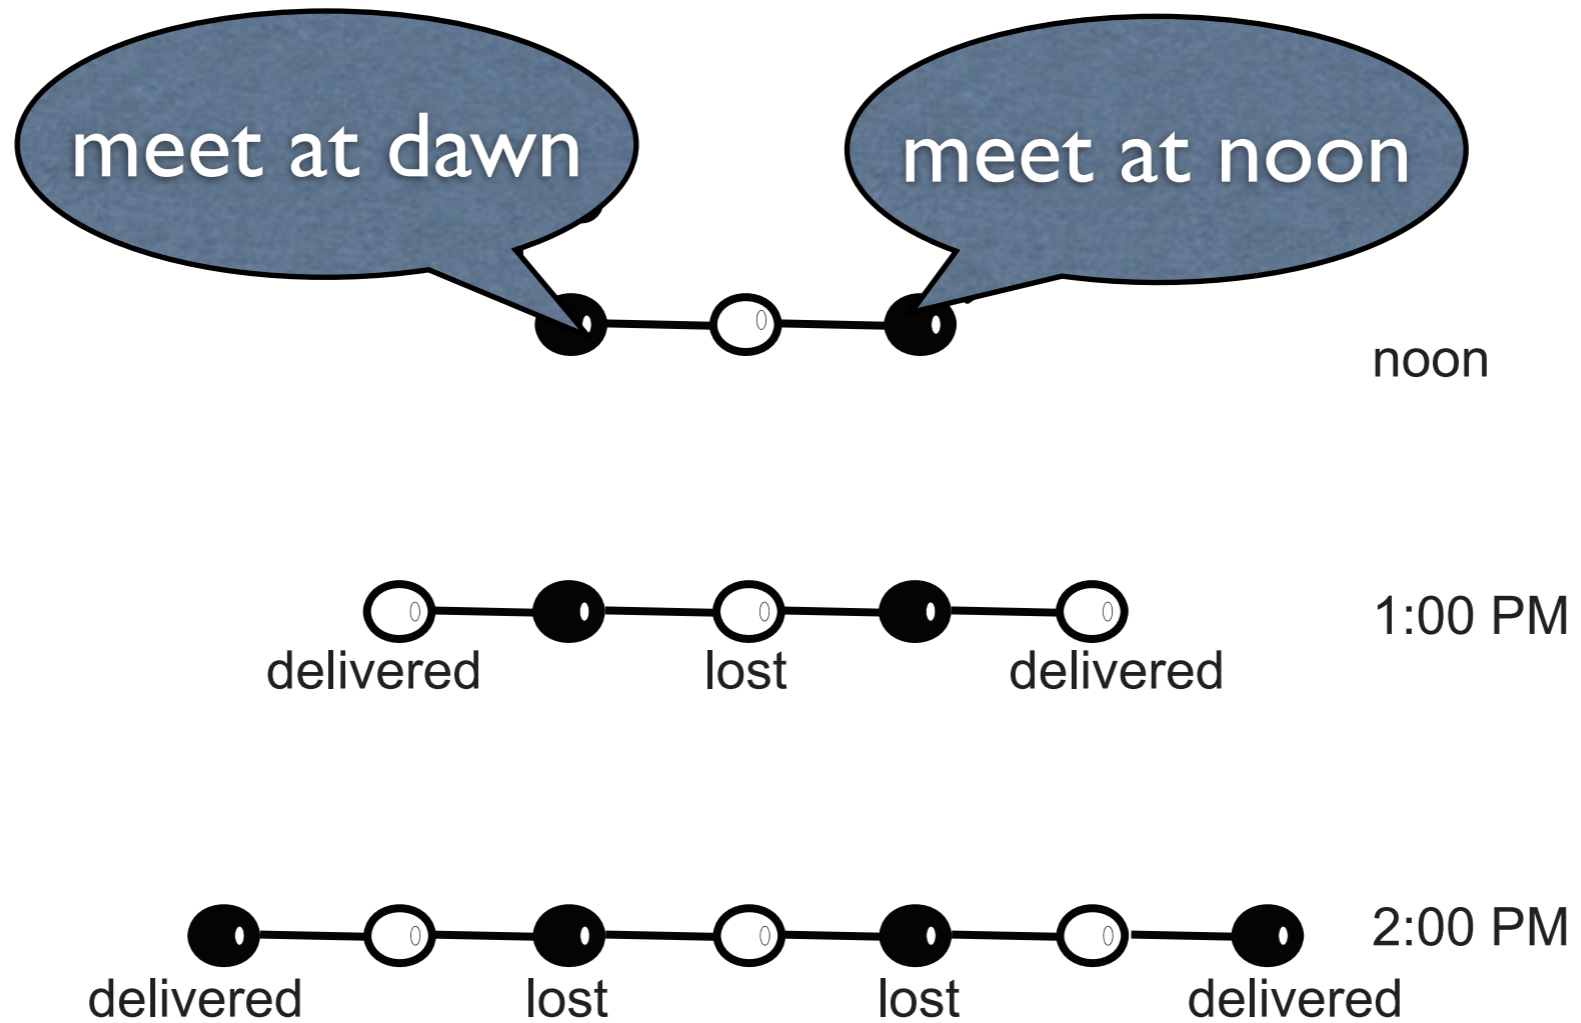
# Combinatorial analysis

- Initially Alice has two possible decisions: meet at dawn, or meet at noon the next day.
- Bob has only one initial state, the white vertex in the middle, waiting to hear Alice's preference.

# Combinatorial analysis

- Initially Alice has two possible decisions: meet at dawn, or meet at noon the next day.
- Bob has only one initial state, the white vertex in the middle, waiting to hear Alice's preference.
- This vertex belongs to two edges (simplexes)


# Evolution




# Topology implies impossibility

No number of successfully delivered acknowledgments will be enough to ensure that show up safely, because the complex is subdivided, and remains connected!

No number of successfully delivered acknowledgments will be enough to ensure that show up safely!



Because not possible to map a connected input complex into a disconnected output complex



**To conclude**



Happy birthday Luis !

