The Asynchronous Computability Theorem

A Marriage Between Distributed Systems and Algebraic Topology

Hei Li

1999	Peter Shor	for Shor's algorithm for factoring numbers in polynomial time on a quantum computer
2000	Moshe Y. Vardi and Pierre Wolper	for work on temporal logic with finite automata
2001	Sanjeev Arora, Uriel Feige, Shafi Goldwasser, Carsten Lund, László Lovász, Rajeev Motwani, Shmuel Safra, Madhu Sudan, and Mario Szegedy	for the PCP theorem and its applications to hardness of approximation
2002	Géraud Sénizergues	for proving that equivalence of deterministic pushdown automata is decidable
2003	Yoav Freund and Robert Schapire	for the AdaBoost algorithm in machine learning
2004	Maurice Herlihy, Michael Saks, Nir Shavit and Fotios Zaharoglou	for applications of topology to the theory of distributed computing
2005	Noga Alon, Yossi Matias and Mario Szegedy	for their foundational contribution to streaming algorithms
2006	Manindra Agrawal, Neeraj Kayal, Nitin Saxena	for the AKS primality test
2007	Alexander Razborov, Steven Rudich	for natural proofs
2008	Daniel Spielman, Shanghua Teng	for smoothed analysis of algorithms
2009	Omer Reingold, Salil Vadhan, Avi Wigderson	for zig-zag product of graphs and undirected connectivity in log space

Two nodes choose between 0 or 1 →
Come to agreement about one of their inputs

Binary consensus problem in asynchronous wait-free model is unsolvable

No timing guarantees in execution

Nodes eventually halt with output value regardless of crashes in other nodes

Hey! Isn't this the FLP result?

Non-generalizable

No insights on properties of model of computation

Inelegant

Given: Number of heads and feet

Find: Number of chickens and rabbits

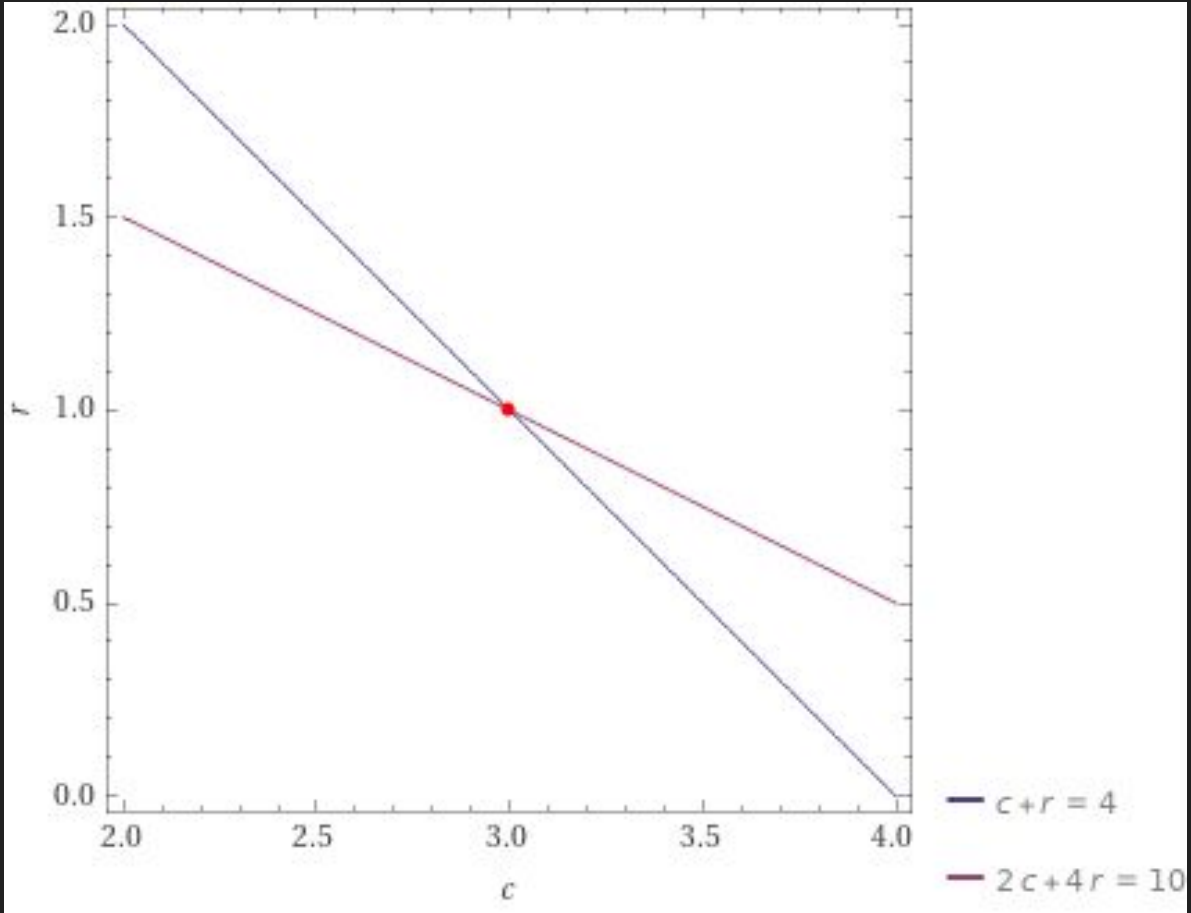

Given: 4 heads, 10 feet

Solution: 3 chickens

+ 1 rabbit

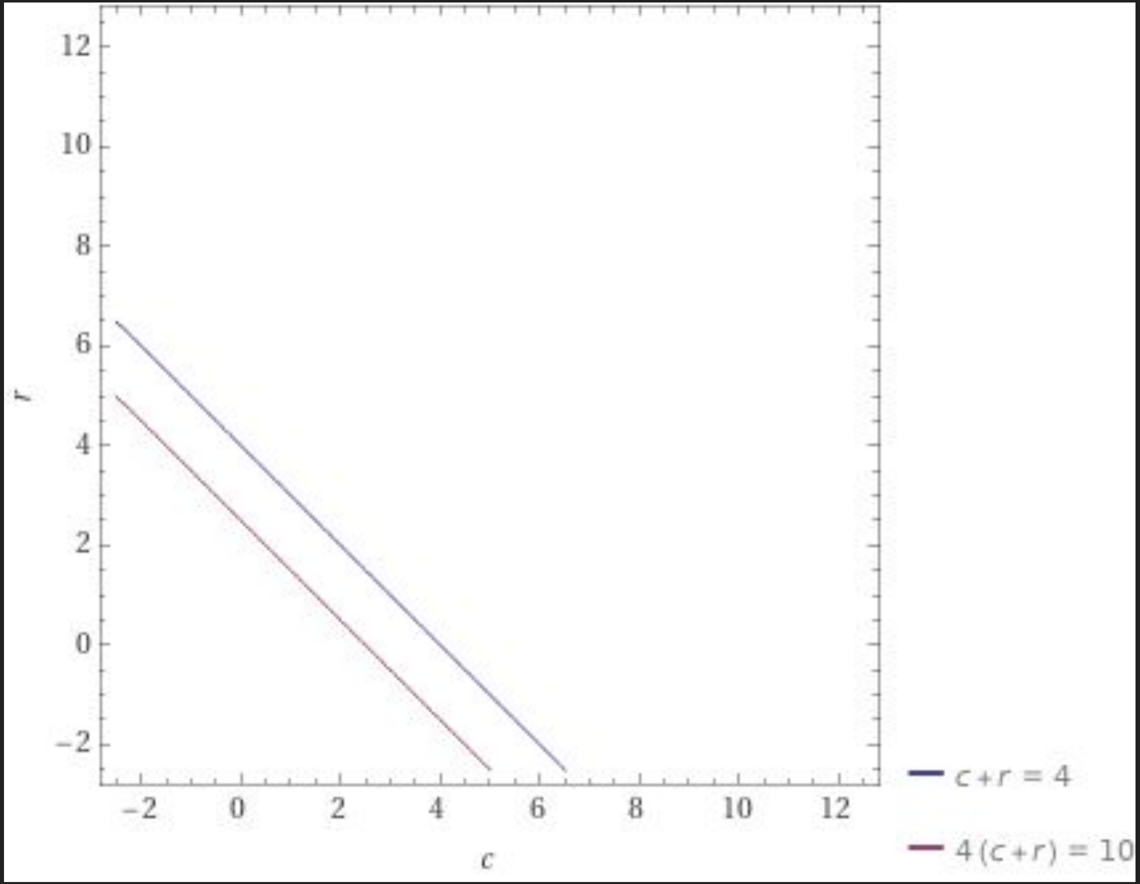

Given: 4 heads, 10 feet

No solutions – parallel lines

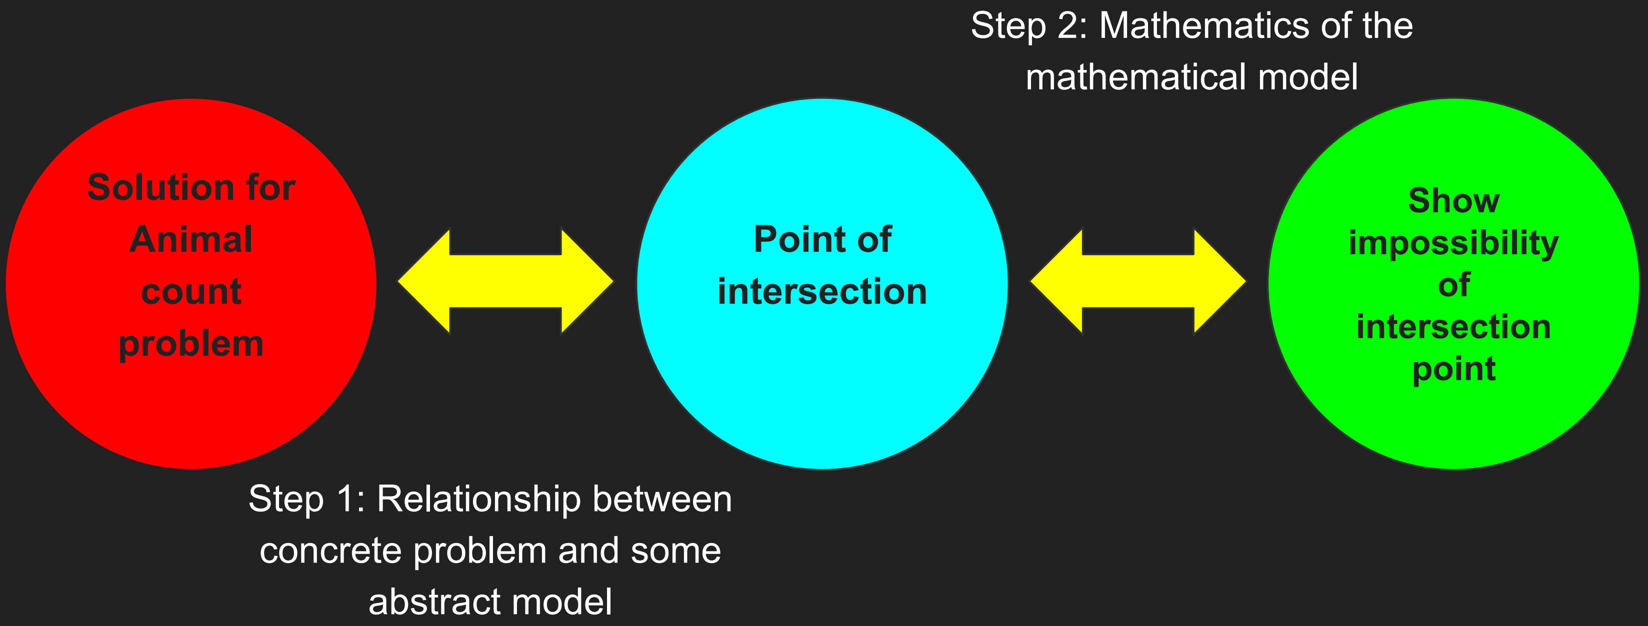
Step 2: Mathematics of the mathematical model
Solution for Animal count problem
Point of intersection
Show impossibility of intersection point
Step 1: Relationship between concrete problem and some abstract model

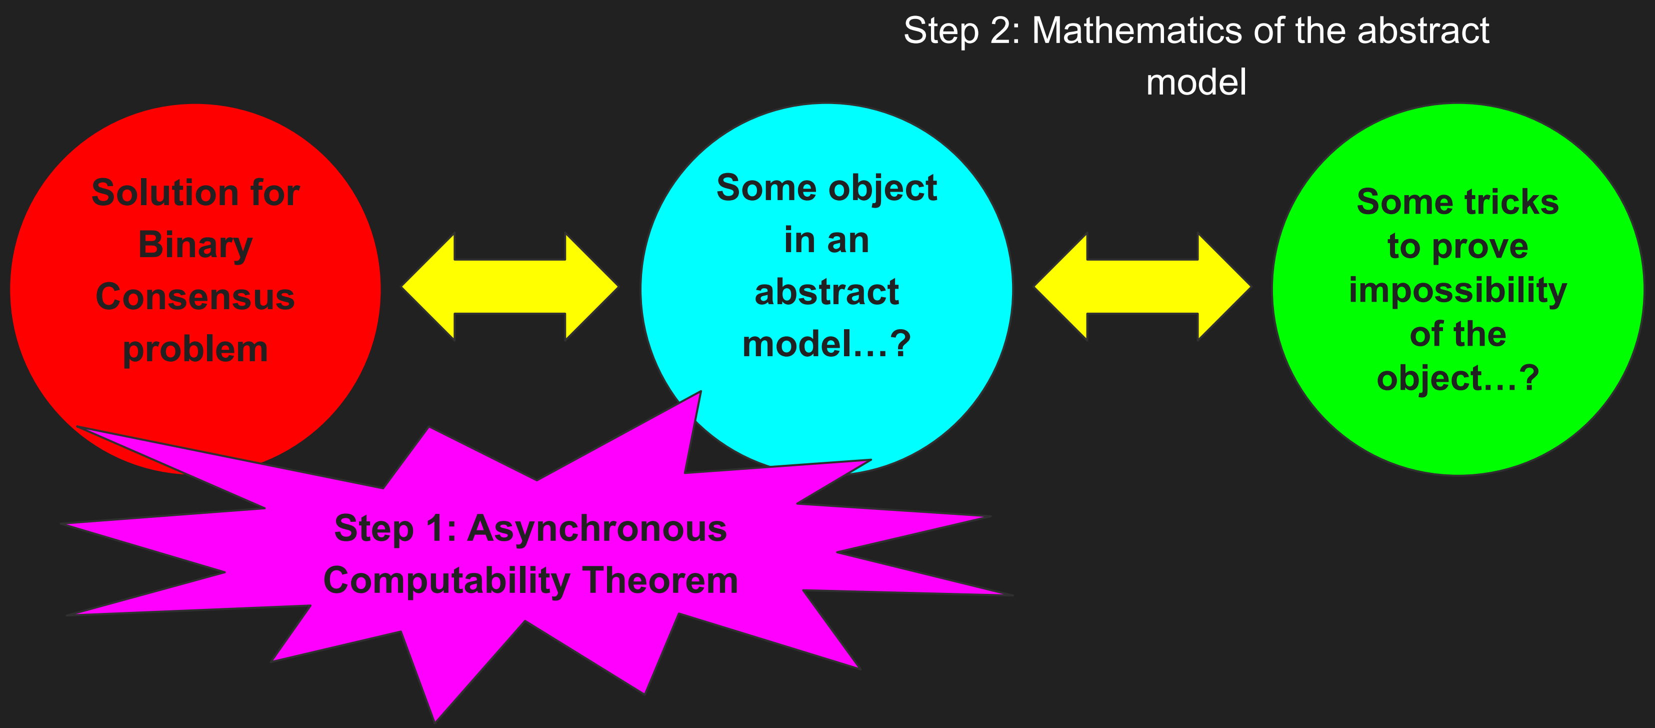
Step 2: Mathematics of the abstract model
Solution for Binary Consensus problem
Some object in an abstract model…?
Some tricks to prove impossibility of the object…?
Step 1: Asynchronous Computability Theorem

A task : <I, O, Δ>

A protocol solves* a task if given any starting input x in I:

Final output is in Δ(x) *

For our binary consensus problem

$I = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$ *

$O = \{(0,0), (1,1)\}$ *

$\Delta((0,1)) = \{(0, 0), (1, 1)\}$

$\Delta((1,1)) = \ldots$

$\ldots\{(1, 1)\}$

A protocol solves* a task if given any starting input x in I:

Final output is in Δ(x) *

We need ***DIMENSION*** to represent clusters of nodes

Simplex is a set of mutually-connected vertices *

Complex is a collection of simplexes *

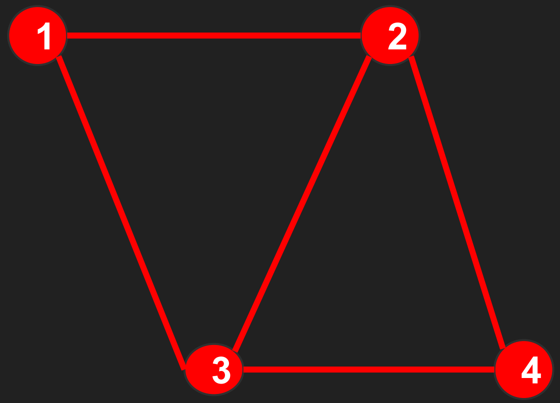

Simplex is a set of mutually-connected vertices *

Complex is a collection of simplexes *

Some simplices: (1, 2, 3), (2, 4), (3),...

The complex (1, 2, 3, 4) is formed by the basic simplices (1, 2, 3) and (2, 3, 4)

(1, 2, 3, 4), (1, 2, 4) are not simplexes!

Given a complex C, a complex σ(C) is a subdivision of C if

- Every simplex in σ(C) is contained in a simplex in C
- Every simplex in C is the union of finitely many simplices in σ(C)

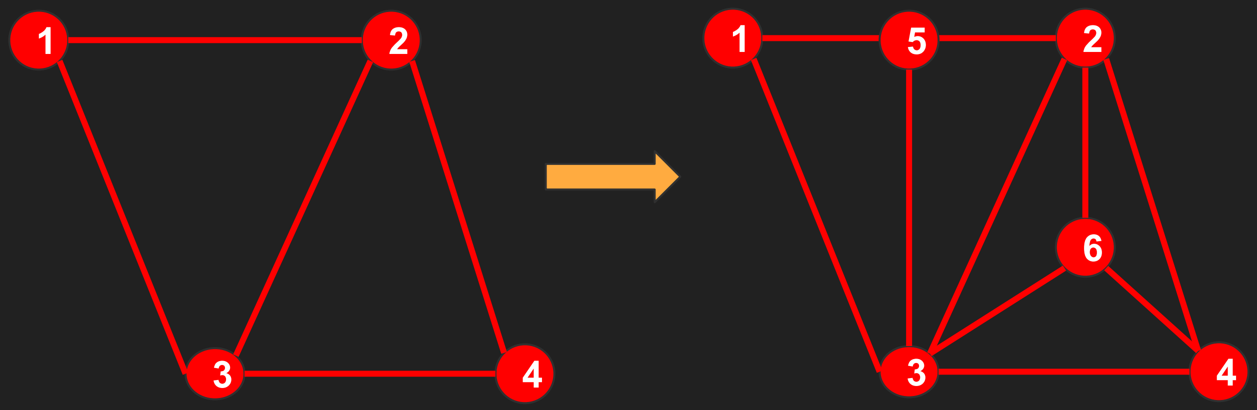

A simplicial map from complex C to complex D, is a function mapping vertices of C to D such that all simplices of C are mapped to simplices of D

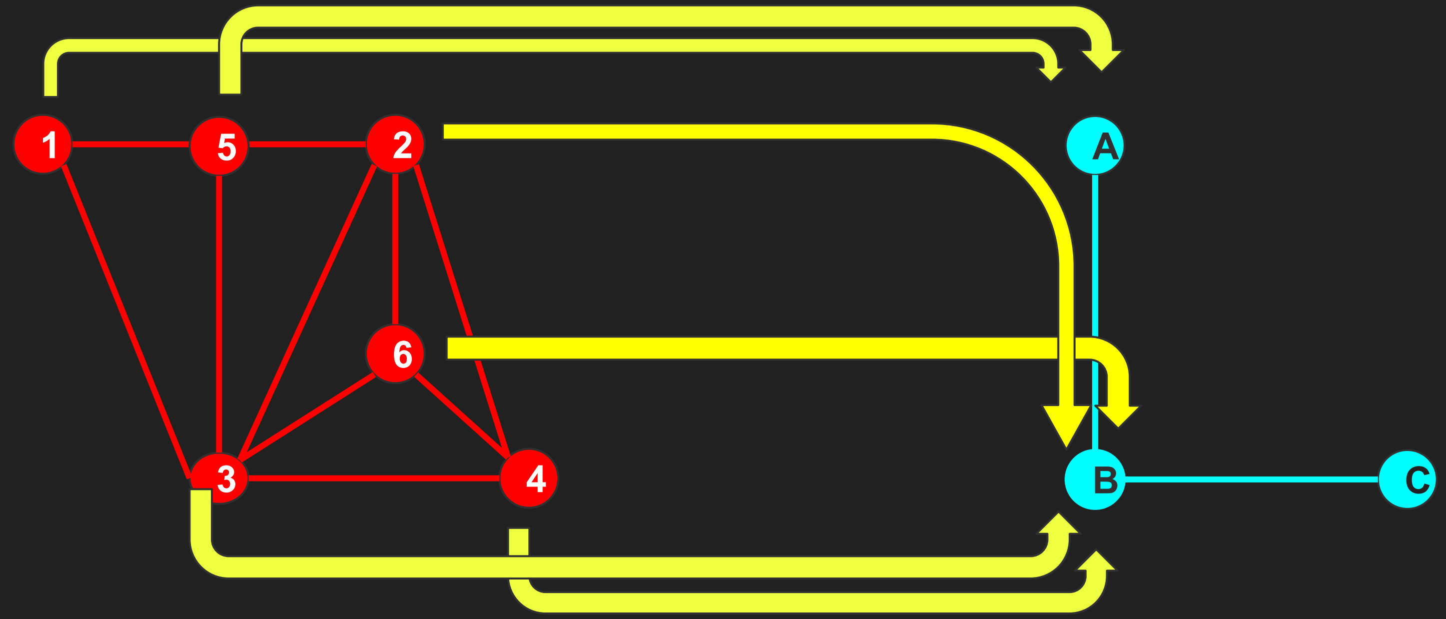

Asynchronous Computability Theorem (Herlihy and Shavit 93)

A decision task <I, O, Δ> has a protocol for an asynchronous wait-free model*

Iff

There exists a subdivision σ of I and a simplicial map μ: σ(I) → O,

such that

It fits Δ requirements *

Stage 2: Decision making

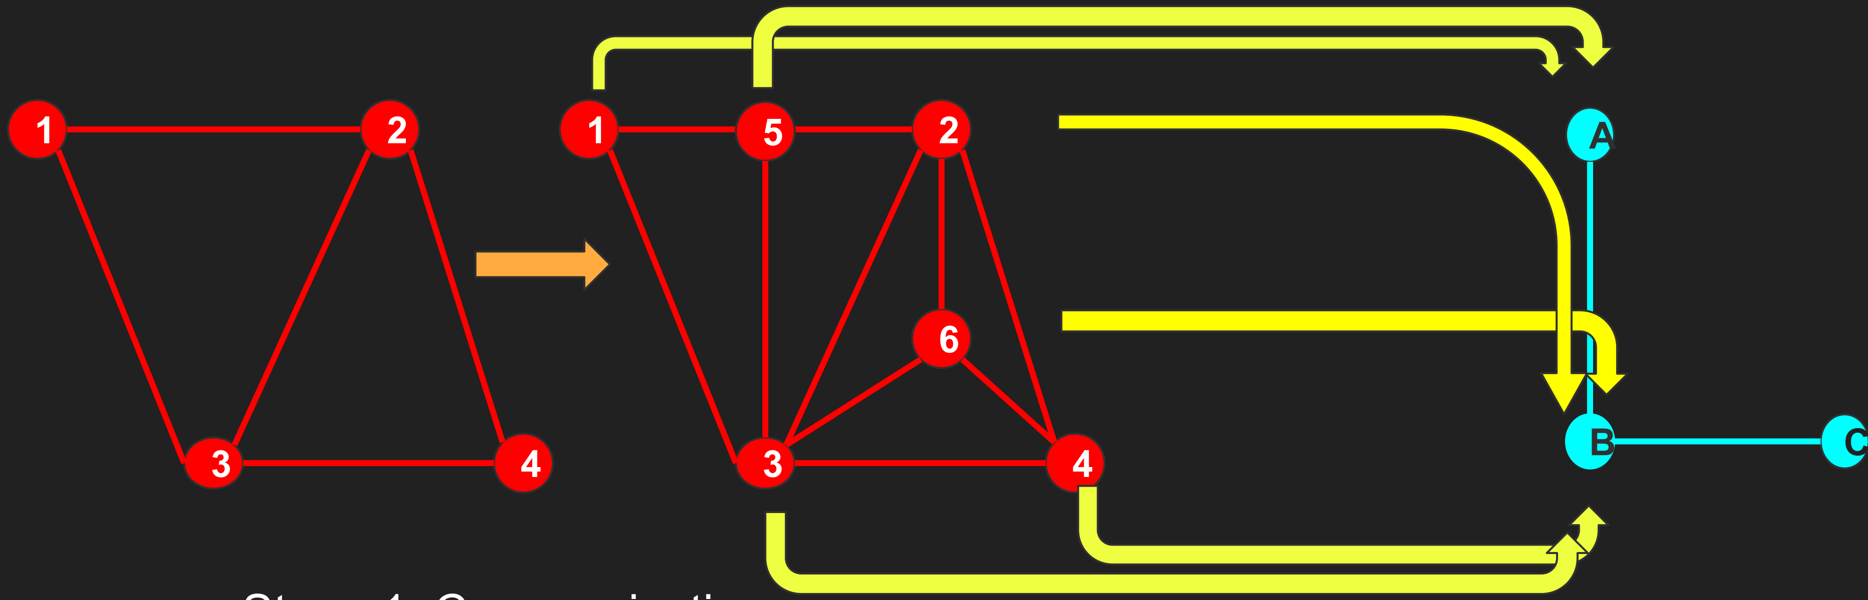

Stage 1: Communication

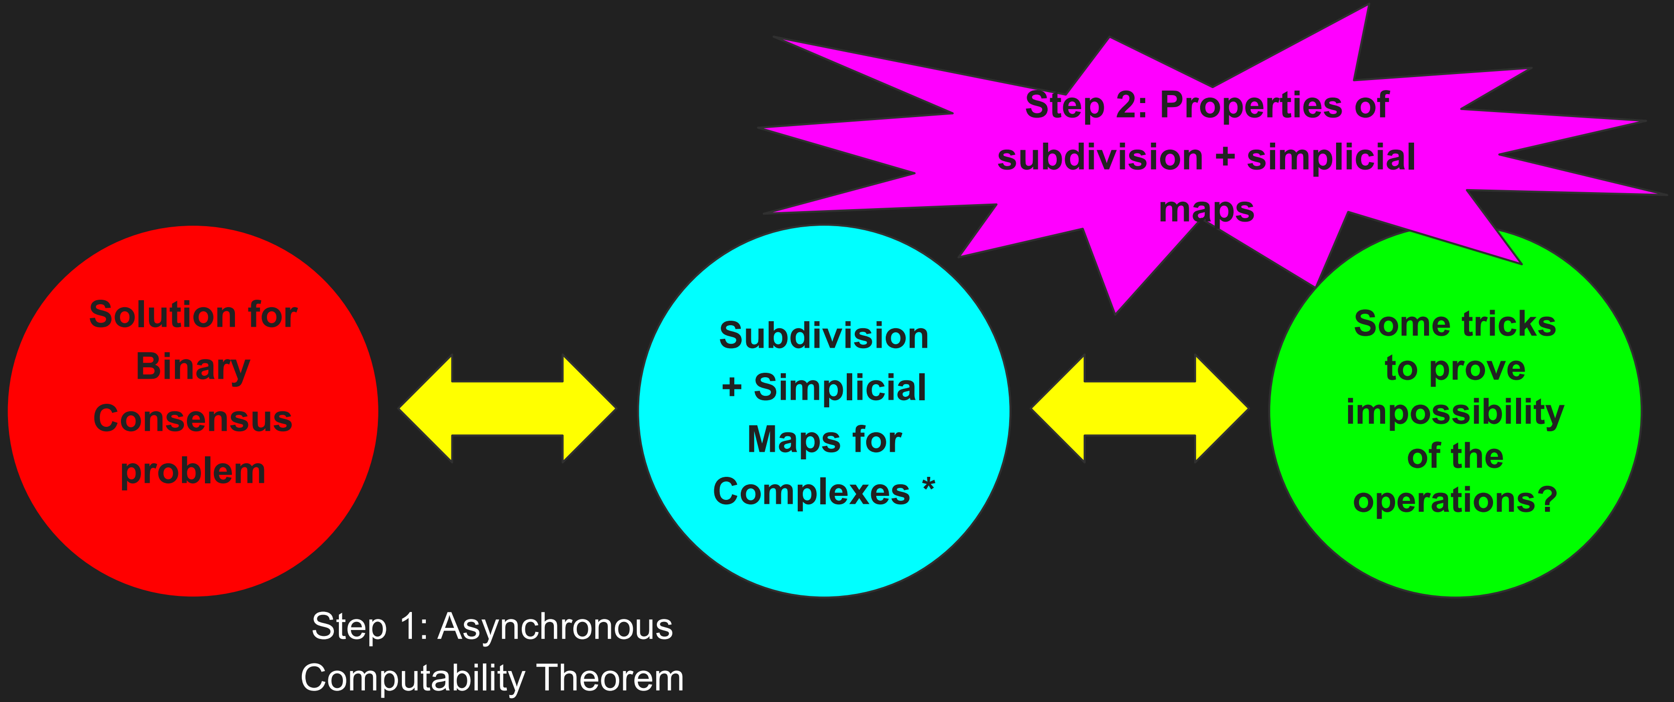

Step 2: Properties of subdivision + simplicial maps
Solution for Binary Consensus problem
Subdivision + Simplicial Maps for Complexes *
Some tricks to prove impossibility of the operations?
Step 1: Asynchronous Computability Theorem

Red = first person

Green = second person

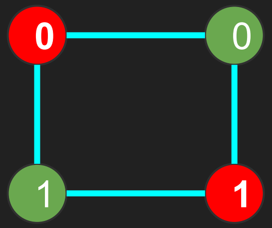

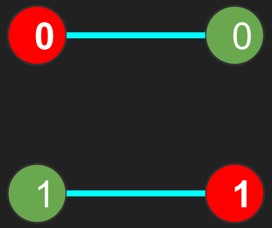

I = {(0, 0), (0, 1), (1, 0), (1, 1)} *

O = {(0,0), (1,1)} *

Consider arbitrary subdivision…

If both parties get 0, they must both terminate with 0.
Red 0 must map back to red 0.

If both parties get 1, they must both terminate with 1.
Green 1 must map back to green 1.

Subdivisions and simplicial maps preserve connectivity!
but red 0 and green 1 are mapped to disconnected components of the output complex

Consider the Quasi-Consensus Problem

Identical to binary consensus problem, but if both are given mixed inputs, either they agree, or green chooses 0 and red chooses 1 (but not vice versa)

After subdivision…

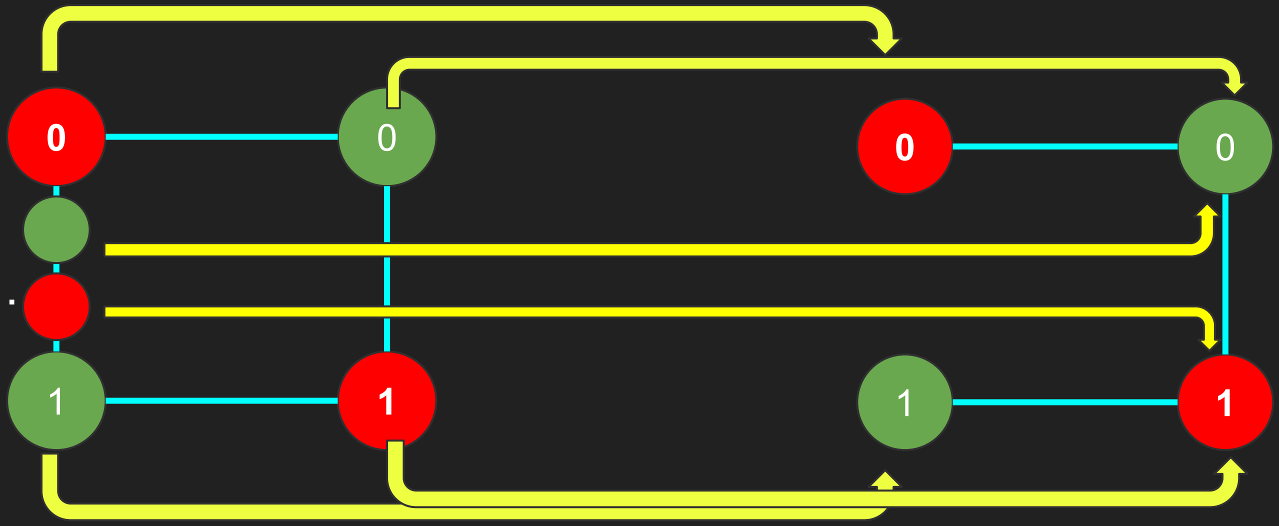

The quasi-consensus problem is solvable!

Consensus problem for more than 2 nodes/Two Generals Problem? Generalization of previous argument

K-set agreement problem for more than 2 nodes? Requires Sperner's Lemma

Anonymous Protocols e.g. renaming problem (Output do not depend on person ID)? Variant of the theorem for anonymous protocols.

Other communication primitives? Herlihy and Rajsbaum 94

Decidability of the protocol? Herlihy and Rajsbaum 97

Complexity of the protocol? Hoest and Shavit 97

1. Protocol for asynchronous wait-free model =
   simplicial map from subdivision of I to O with certain Δ properties *

2. Binary consensus problem cannot be solved since one cannot construct subdivision + simplicial map due to connectivity property of the map

3. **Topological perspective for theory of distributed and concurrent computation (or other branches of computer science…?)**

Thank you

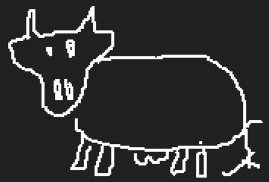

Colouring

1. A complex is chromatic if
   - Each vertex has a colour and no “adjacent” vertices have the same colour
2. A simplicial map is chromatic if
   - It also preserves vertex’s colours after the map

Intuition: the colour represents a single person in the protocol

Other handwavy definitions

A carrier is the unique smallest simplex in the original complex that contains that simplex in the subdivision complex

A subdivision is chromatic if it is a chromatic complex and for each simplex S in the subdivision, the colours of S is in the set of colours of carrier S

What is a good mathematical model
for a distributed systems task?

Turns out using a graph is not good enough!
We need ***DIMENSION*** to represent clusters of nodes as well.

Simplex is a non-empty finite set

Complex is a collection of simplices closed under containment

- Any subset of a simplex in a complex is also a simplex of the complex

Asynchronous Computability Theorem (Herlihy and Shavit 93)

A decision task (I, O, Δ) has a wait-free protocol using read-write memory

Iff

There exists a subdivision σ of I
and a simplicial map μ: σ(I) → O,

such that

σ is a chromatic sub-division and
μ is chromatic simplicial map and
for each simplex S in σ(I), μ(S) ∈ Δ(carrier (S, I))

Our input complex I looks like this:

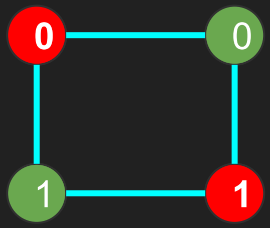

Our output complex O looks like this:

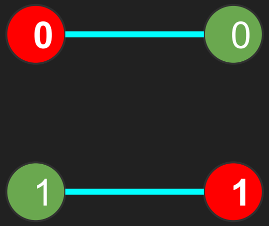

Can you guess why we cannot find the required subdivision of the input complex + simplicial map required?

Similarly:

Simplex = state of a group of people

Subdivision = possible states after running a protocol

Common vertex in two simplexes= person who cannot distinguish two states based on their local information

Simplicial map that fits Δ= what each person chooses after running the protocol based on their local state

Simplex = set of mutually connected nodes

Complex = collection of simplices

Subdivision = triangulation of a complex

Simplicial map =
mapping vertices of one complex to another while preserving simplexes