Topology and Distributed Computing

Maurice Herlihy
Brown University
Nir Shavit
Sun Microsystems & Tel Aviv U.

From the New York Times …

SAN FRANCISCO, May 7 - Intel said on Friday that it was scrapping its development of two microprocessors, a move that is a shift in the company's business strategy….

The NYT Continues …

Intel … [has] decided to focus its development efforts on “dual core” processors …. with two engines instead of one, allowing for greater efficiency because the processor workload is essentially shared.

NYT 8 May 2004

The Future of Computing

- Speeding up uniprocessors is harder and harder
- Intel, Sun, AMD now focusing on "multi-core" architectures
- Soon, every computer will be a multiprocessor

What does this have to do with Theory?

Alan Mathison Turing.

Inventor of the Turing Machine.

Turing Computability

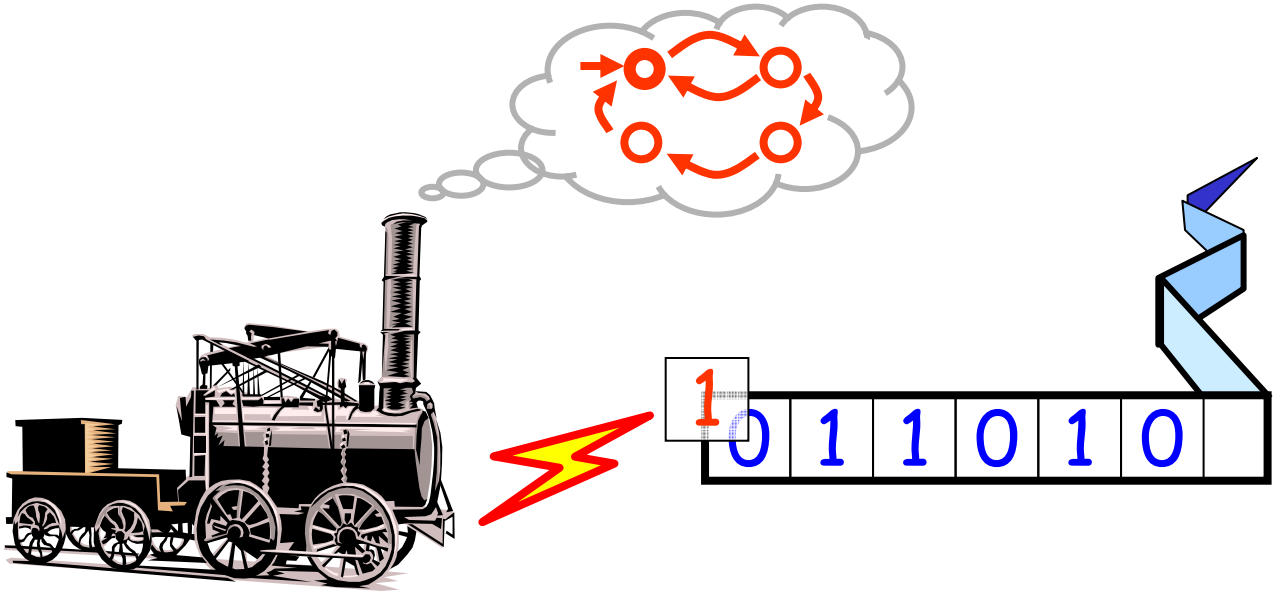

- Mathematical model of computation
- What is (and is not) computable on uniprocessors

Time and Asynchrony

"Time is Nature's way of making sure that everything doesn't happen all at once." (Anonymous, circa 1970)

Real world asynchrony: no clock in the sky

Asynchronous Computability?

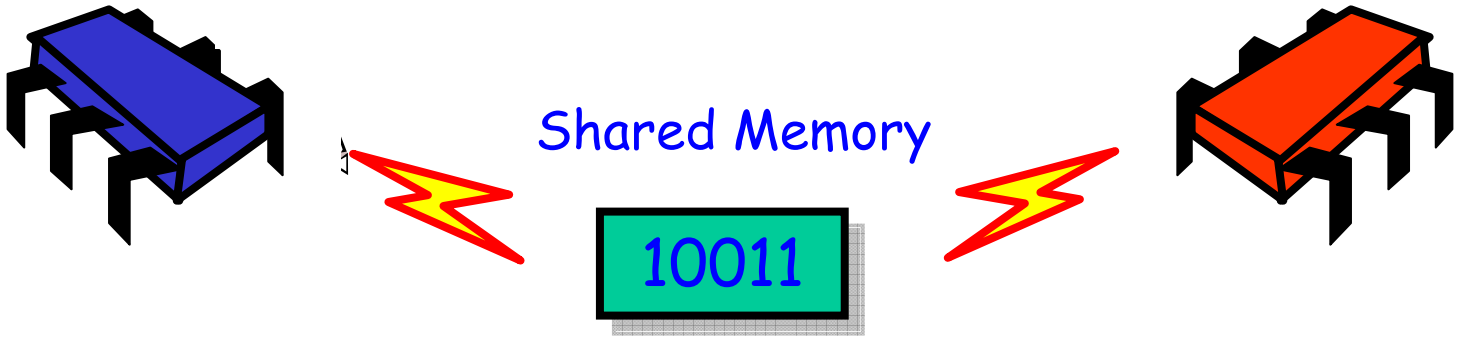

- Mathematical model of asynchronous computation
- What is (and is not) asynchronously computable

Asynchrony Complicates Life

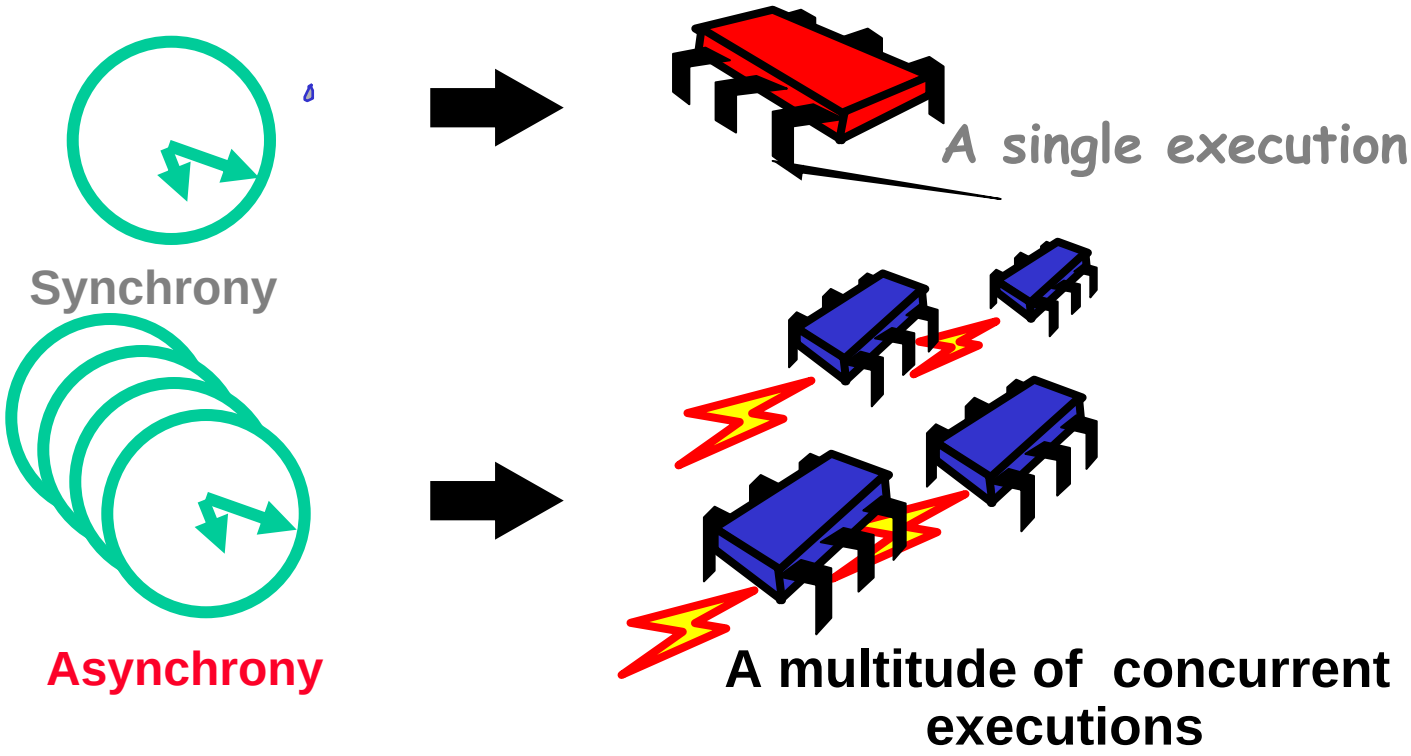

Seek help from modern mathematics!

FLP

- Fischer Lynch Paterson
- Showed that asynchronous computability ≠ Turing computability
- Consensus
  - Trivial in uniprocessor
  - Impossible with 1 asynchronous failure
- Reasoned directly about executions

Graph Theory

- Biran, Moran, Zachs 1988
- Single asynchronous failure
- Coordination problem
  - is a graph
- Problem is asynchronously computable
  - iff graph is connected

More Graph Theory

- Fischer Lynch 82, Dolev Strong 83, Merritt 85, Dwork Moses 90, …
- Synchronous crash failures
- Computation state
  - is a graph
- Consensus not asynchronously computable
  - while graph is connected

Limitations of Graph Theory

- Asynchronous model
  - Multiple failures?
- Synchronous model
  - Problems beyond consensus?
- Need a more general notion of connectivity

Topological Approach

- Borowsky Gafni STOC 93
- Herlihy Shavit STOC 93, JACM 1999
- Saks Zaharoglou STOC 93, SIAM 2000
- Showed k-set agreement impossible
  - Generalization of consensus [Chaudhuri 90]
  - Open problem for several years

The Topological Approach

- Computations as geometric object
  - Use topological methods to show existence of "bad" execution.
- Borowsky & Gafni
  - Sperner's Lemma
- Herlihy & Shavit
  - Simplicial complexes & homology
- Saks & Zaharoglou
  - Brouwer fixed-point theorem

Example: Autonomous Air Traffic Control

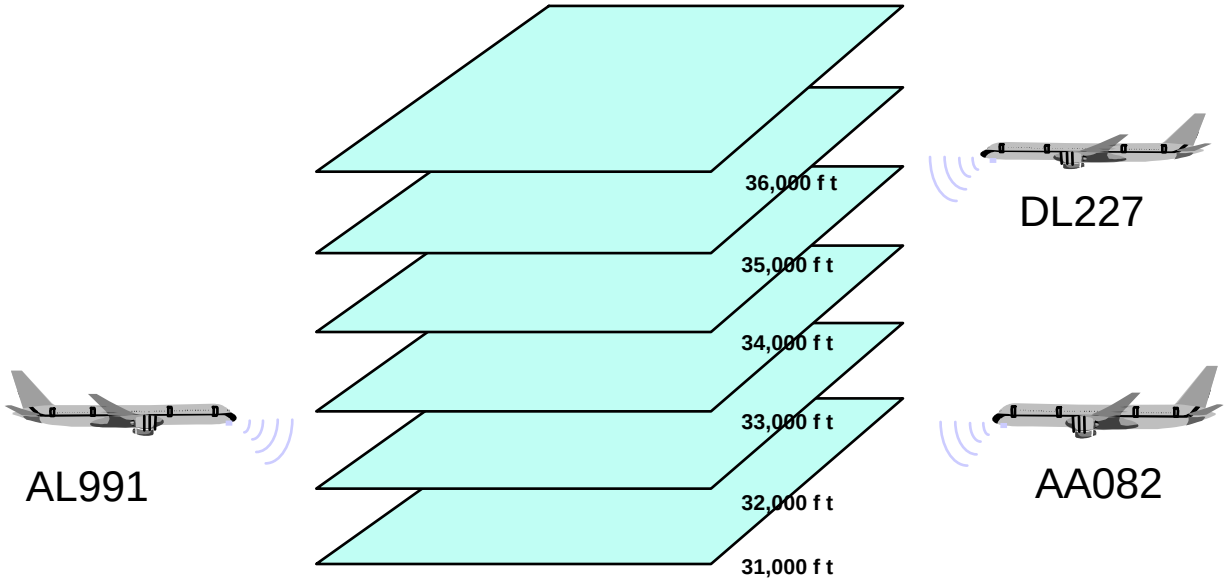

Pick your own altitude.

How many slots do we need to allow safe coordination?

Example Task: Renaming

- Process has input name (flight #)
- Must generate output name (altitude)
- Interested in comparison-based protocols:
  - Equality: A=B?
  - Order: A<B?
  - Nothing else (rules out trivial solutions)

History

- Proposed by Attiya, Bar-Noy, Dolev, Peleg, , Reischuk
- They showed
  - Solution for 2n+1 names
  - impossibility for n+2 names
  - Intermediate values hot open question
- Long-standing open problem …
  - Topological methods showed …
  - Intermediate values also impossible

A Vertex

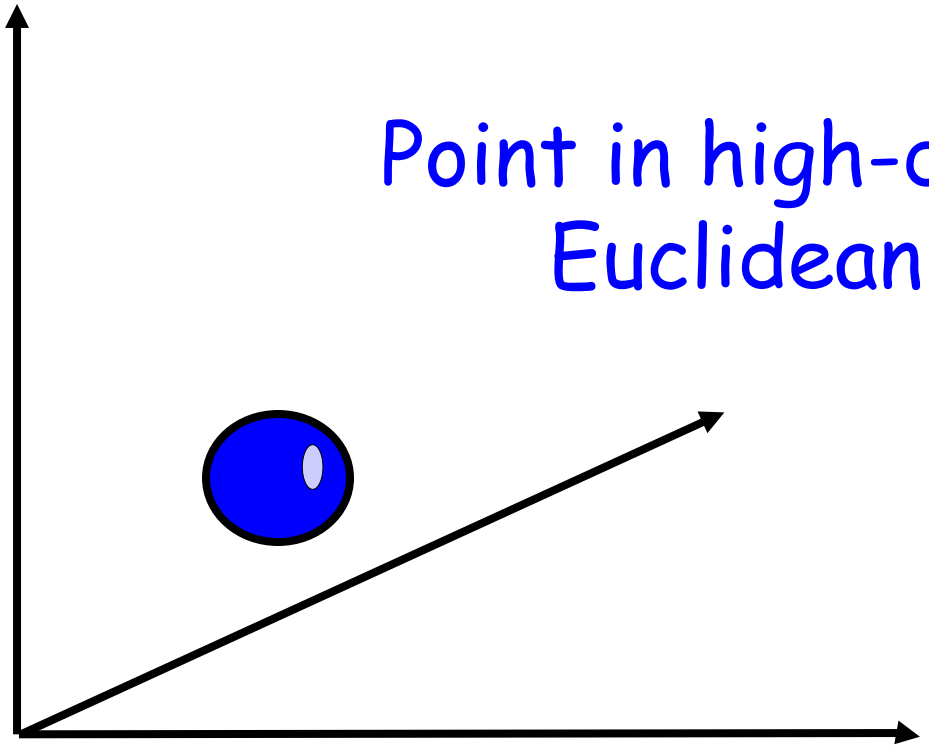

Point in high-dimensional Euclidean Space

Simplexes

0-simplex
(vertex)

1-simplex
(edge)

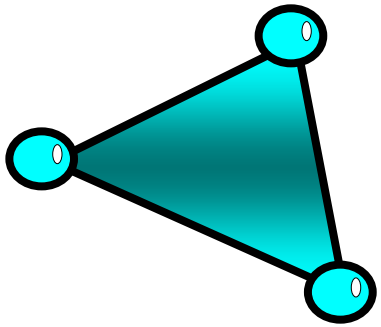

2-simplex
(solid triangle)

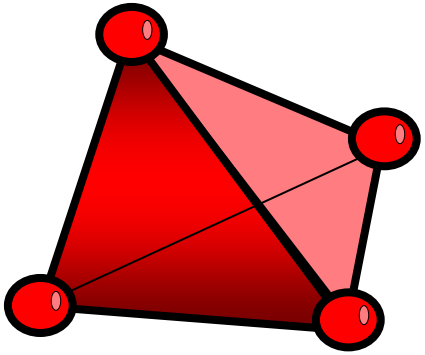

3-simplex
(solid tetrahedron)

Simplicial Complex

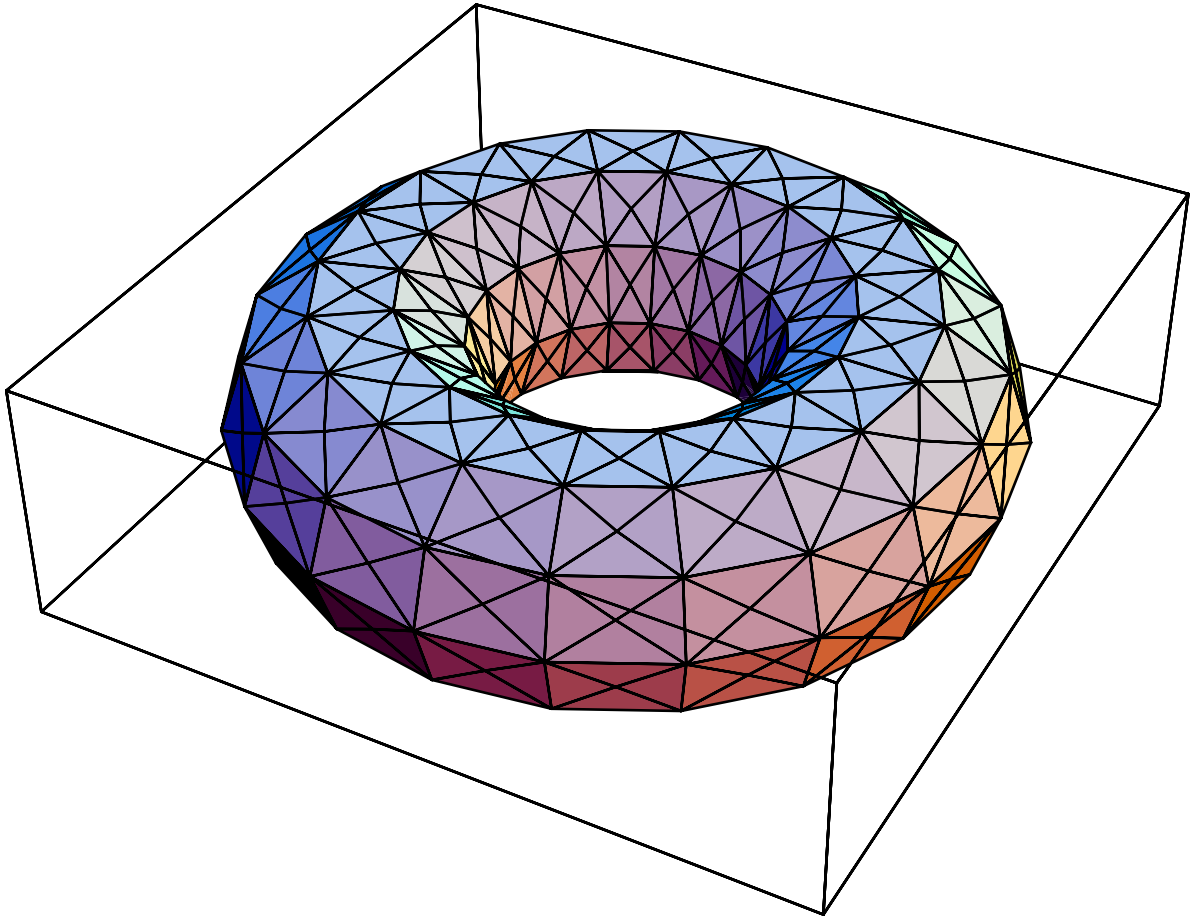

A Cycle

OK to think of it as oriented path

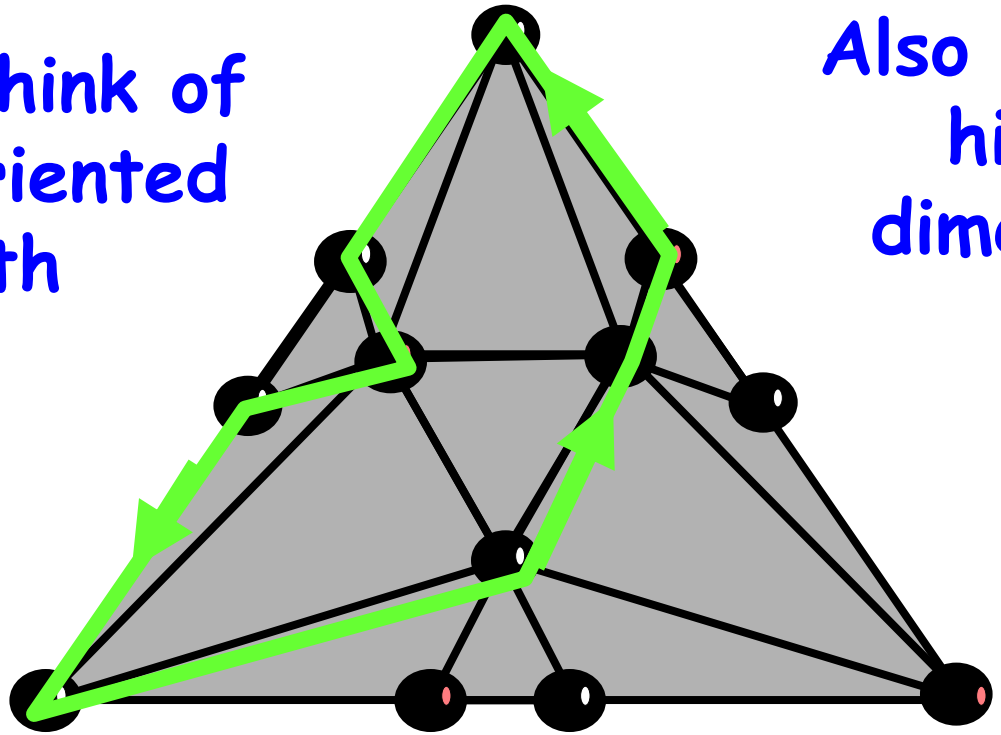

Also works in higher dimensions

A Boundary

A cycle is a boundary if it goes around a "solid" region

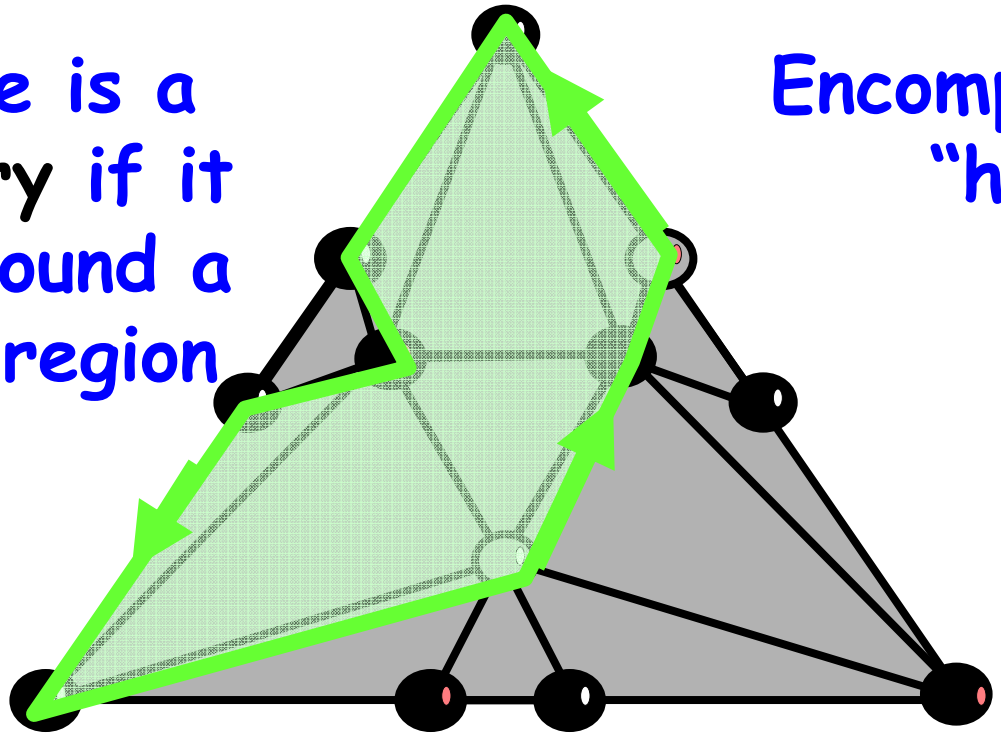

Encompasses no "holes"

Not Every Cycle is a Boundary

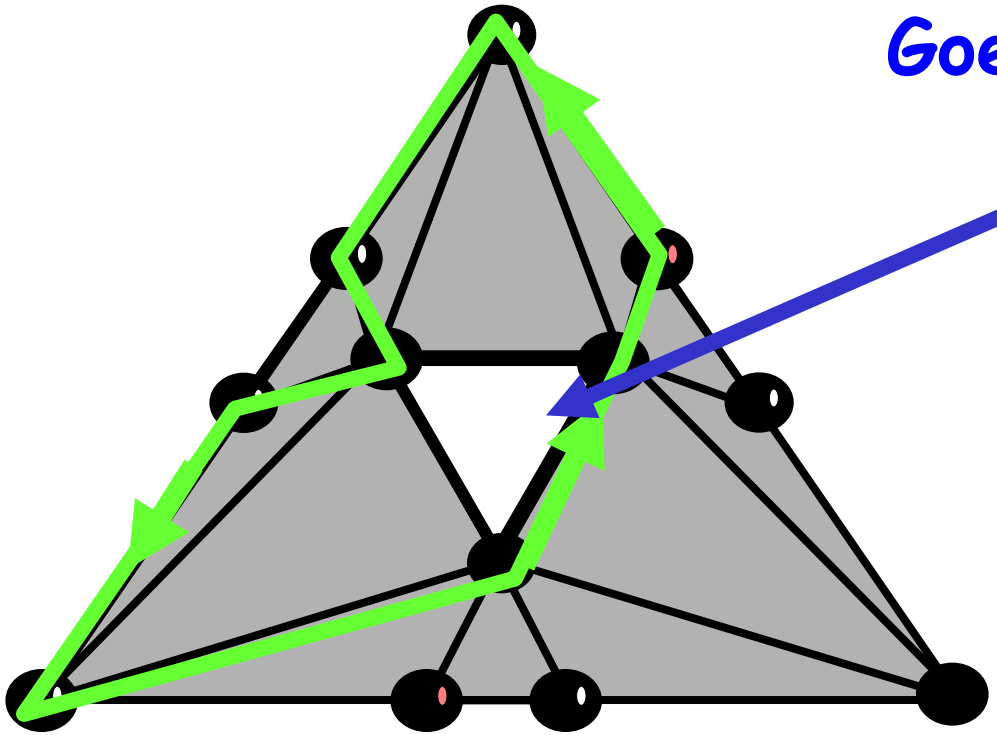

Connectivity

- A complex is n-connected if
  - Every cycle of dimension n or less
  - Is also a boundary
- No "holes" in any dimension
  - Fundamental group is trivial
  - Higher homology groups trivial

Chromatic Complexes

Each n-simplex colored by n distinct colors

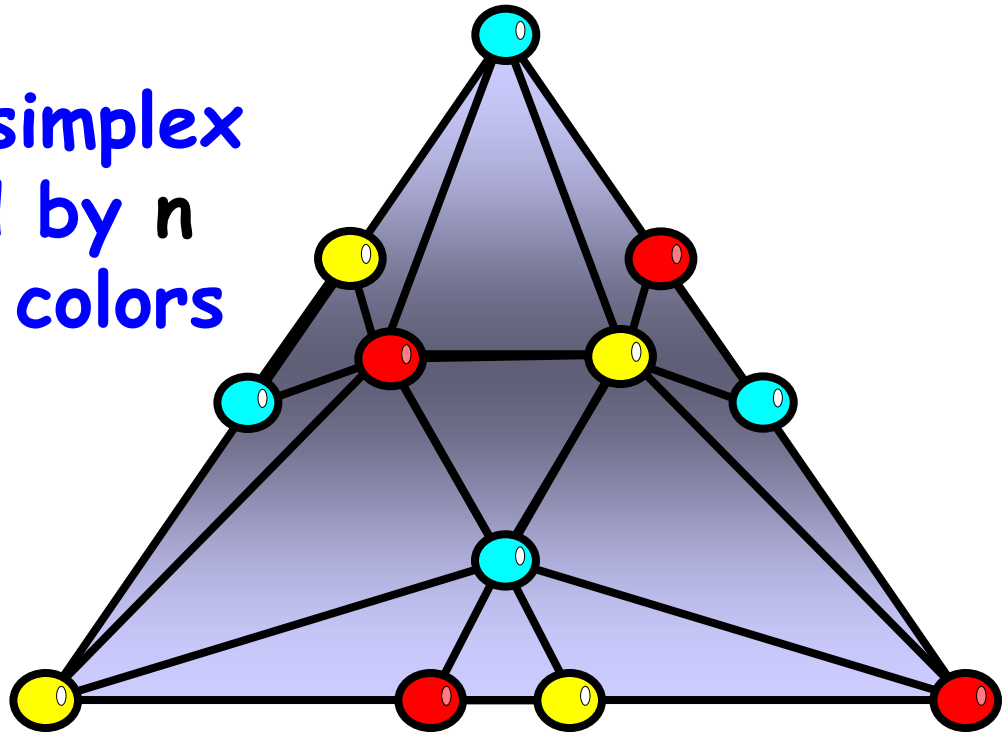

Corresponding to process ids

Chromatic Simplicial Map

Color-preserving vertex-to-vertex map

That also carries simplexes to simplexes

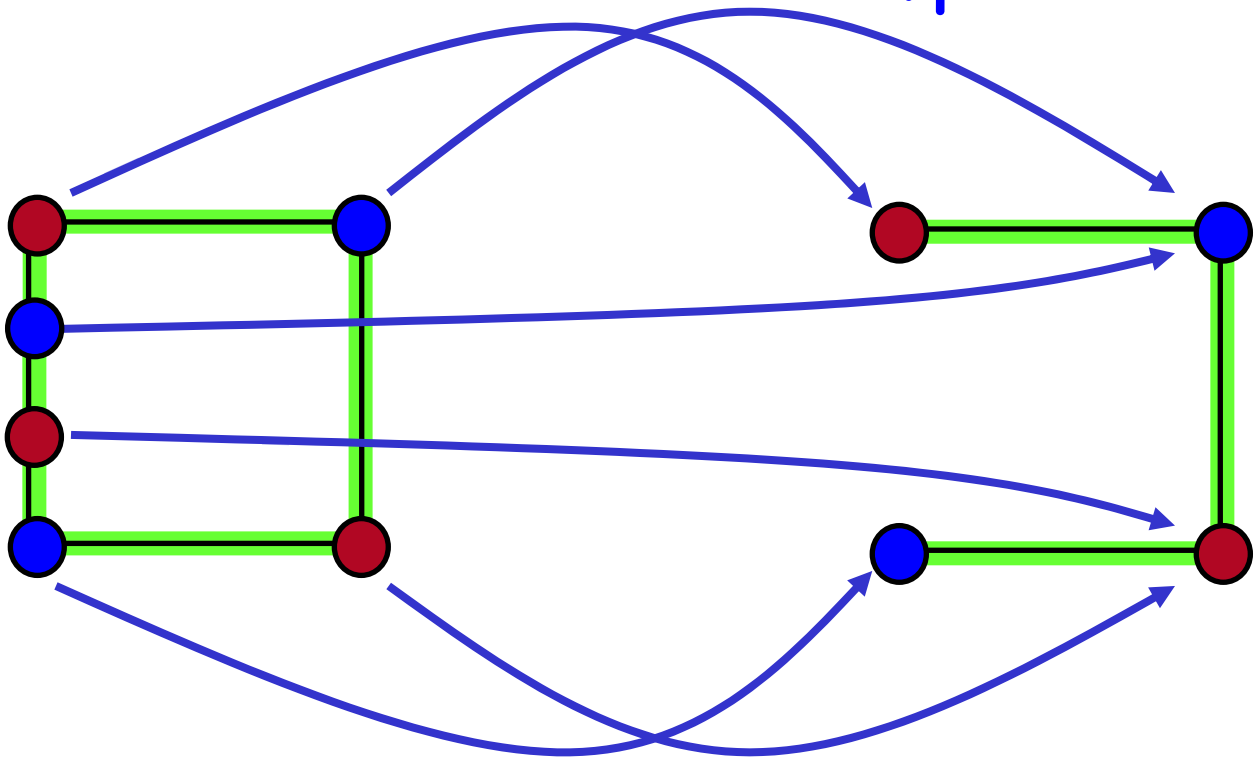

Chromatic Simplicial Maps Preserve Boundaries

A simplicial map can't wrap a boundary around a "hole"

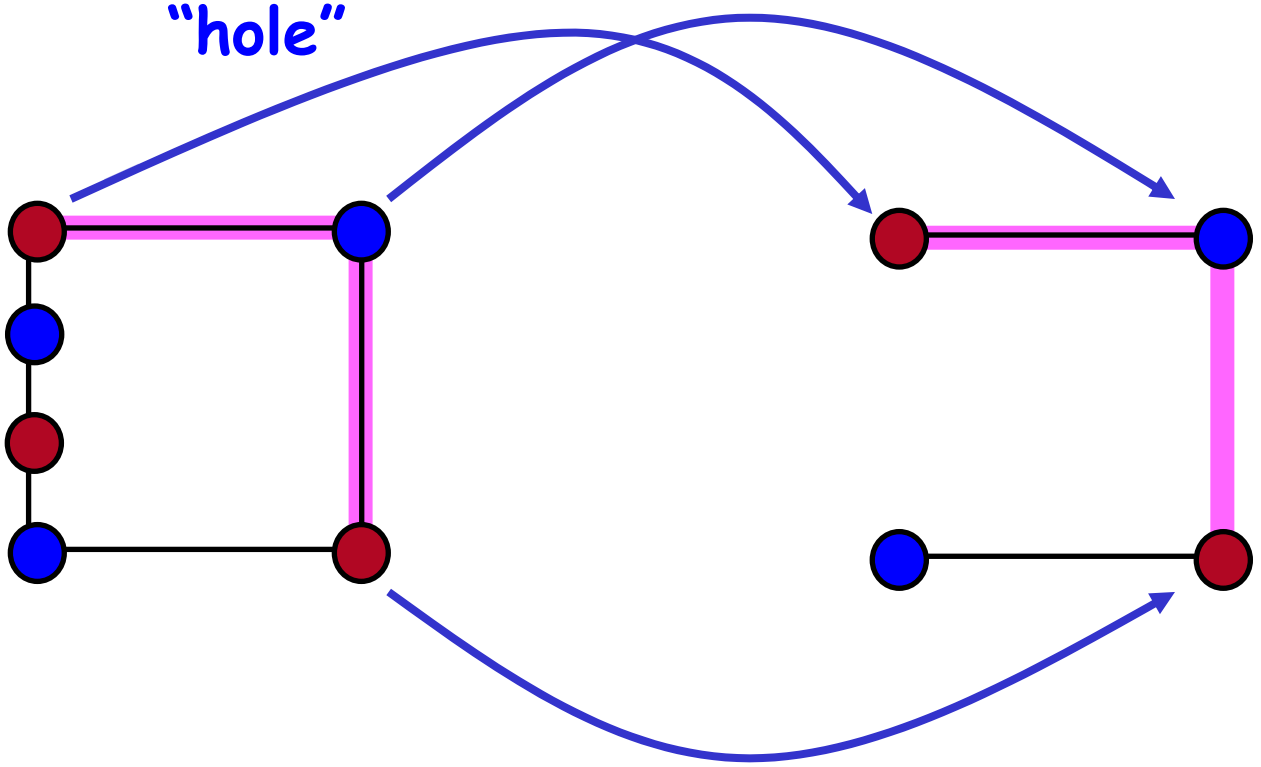

Vertex = Process State

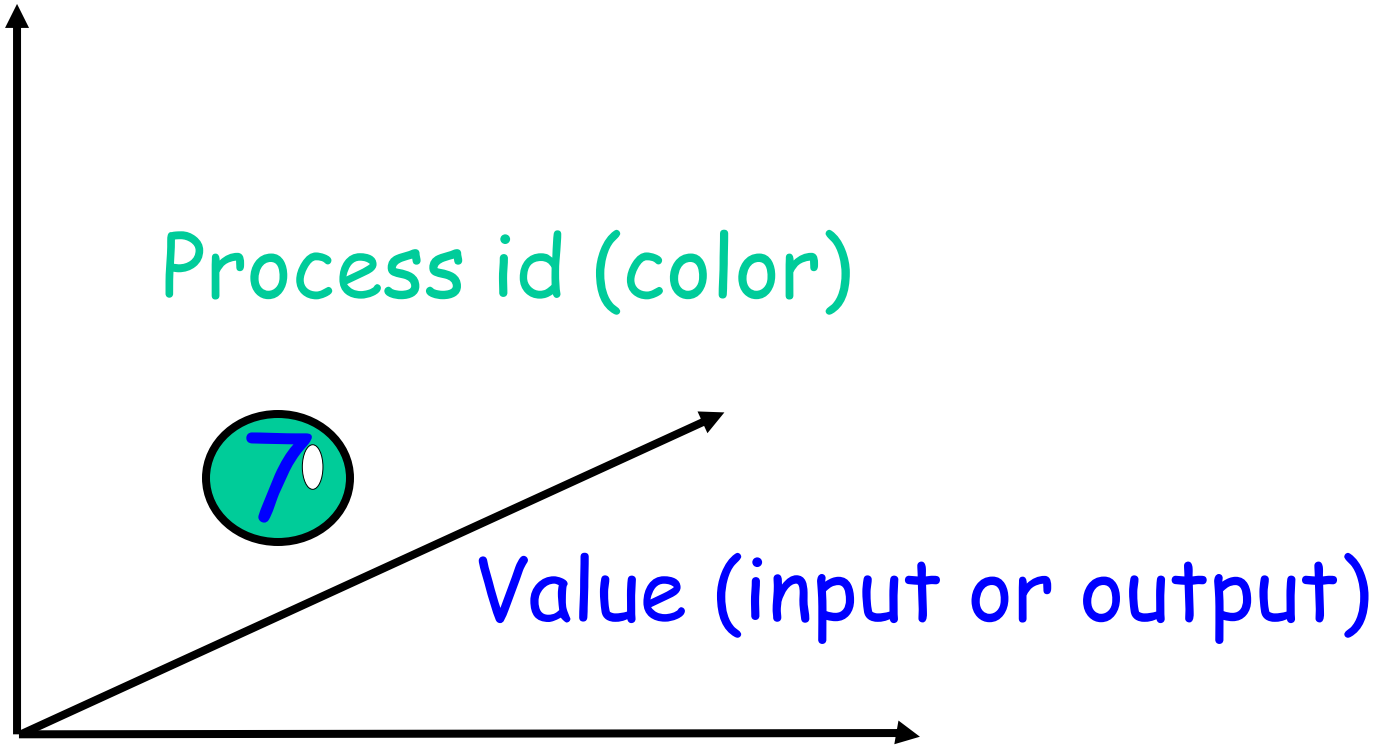

Simplex = Global State

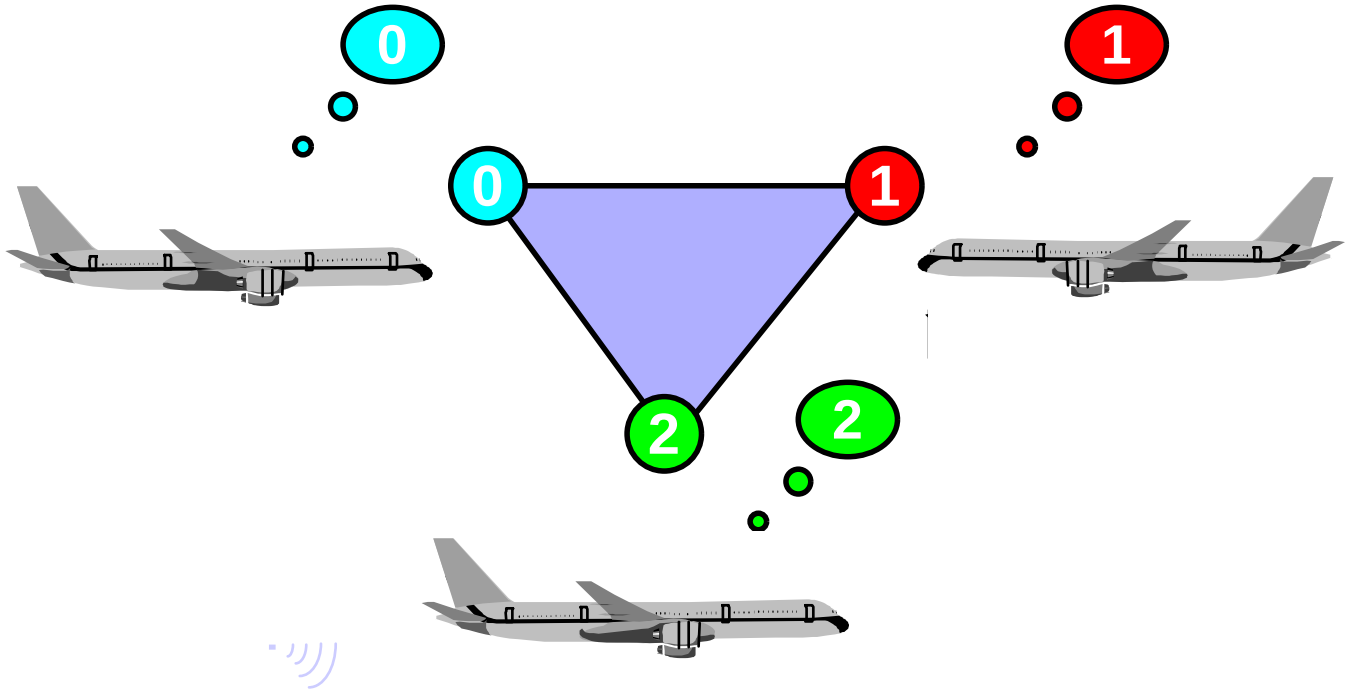

Complex = Global States

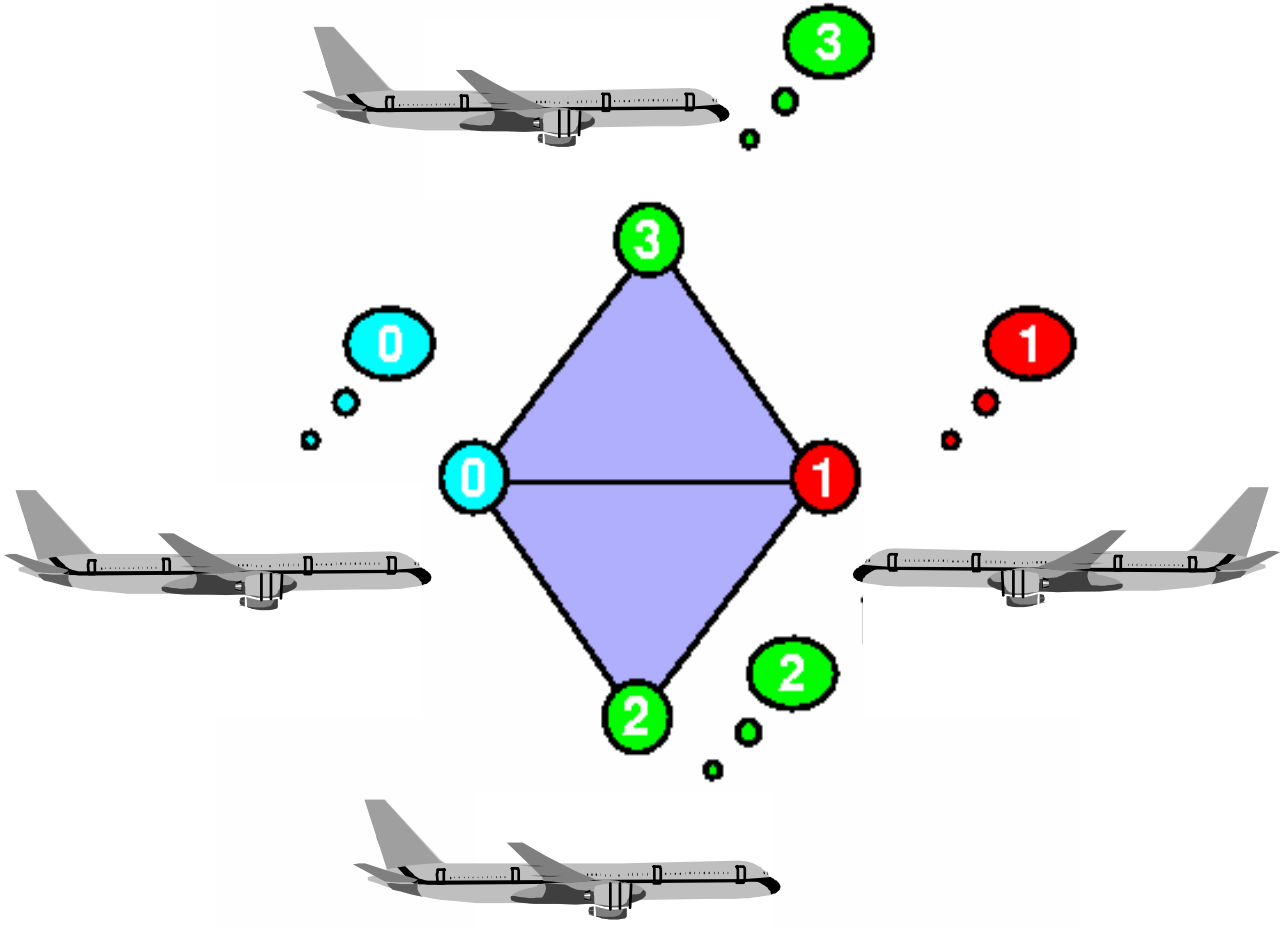

Decision Task

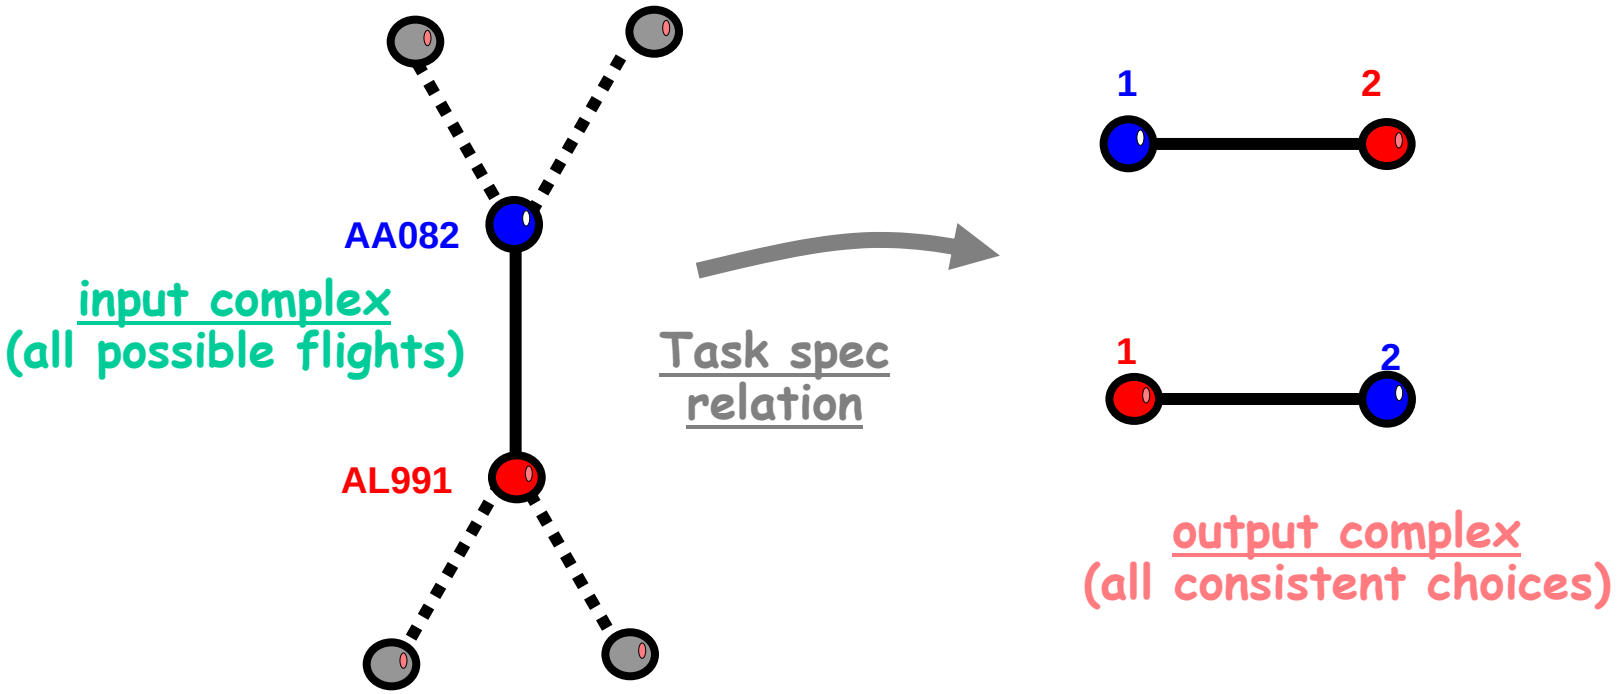

3 Planes, 3 Slots

input complex **output complex**

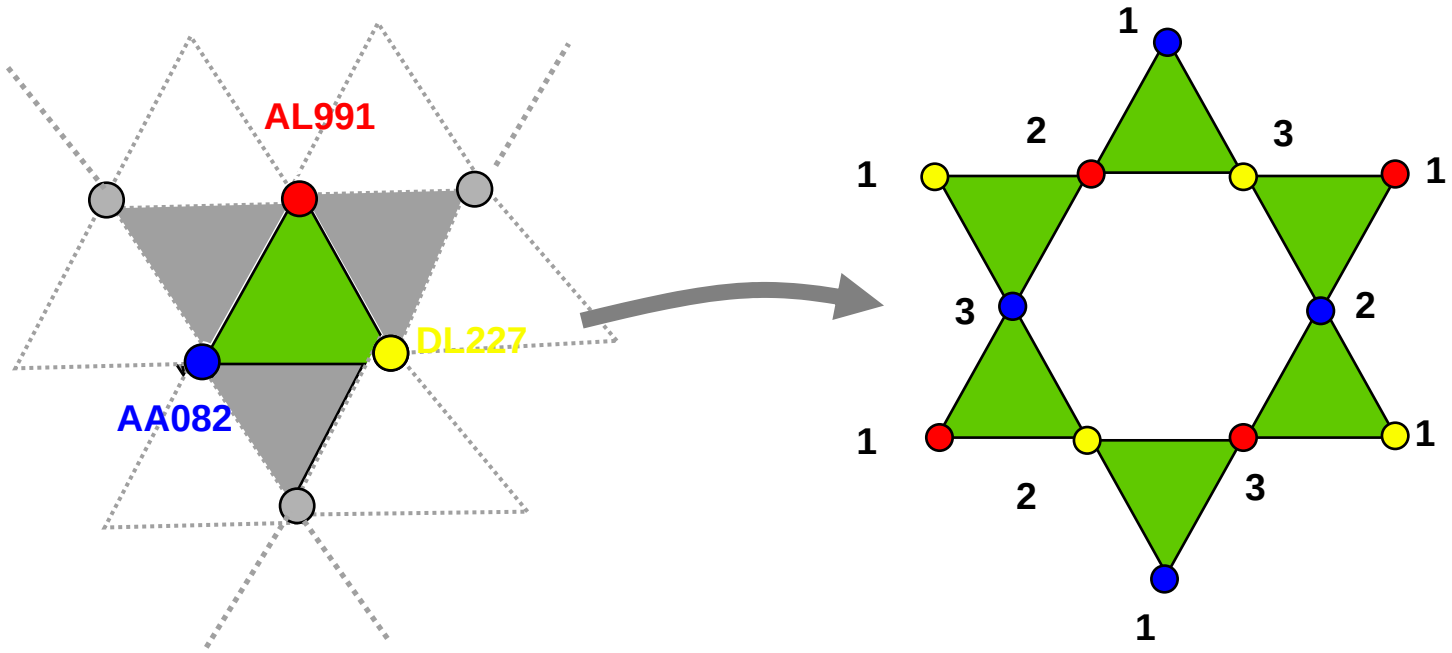

input complex **output complex**

3 Planes, 4 Slots

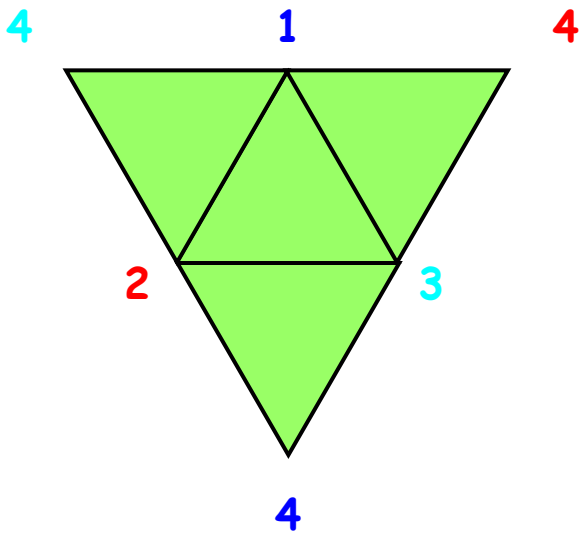

3 Planes, 4 Slots

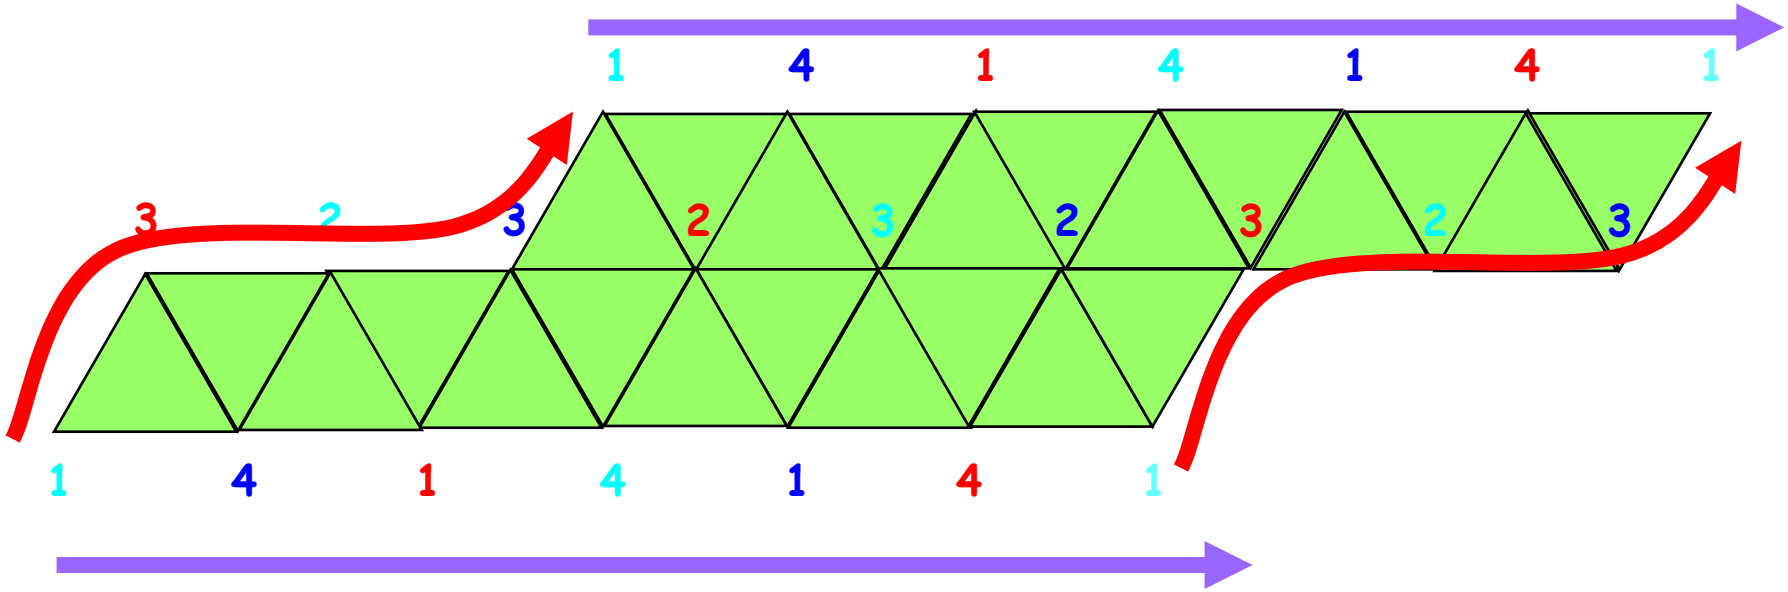

3 Planes, 4 Slots

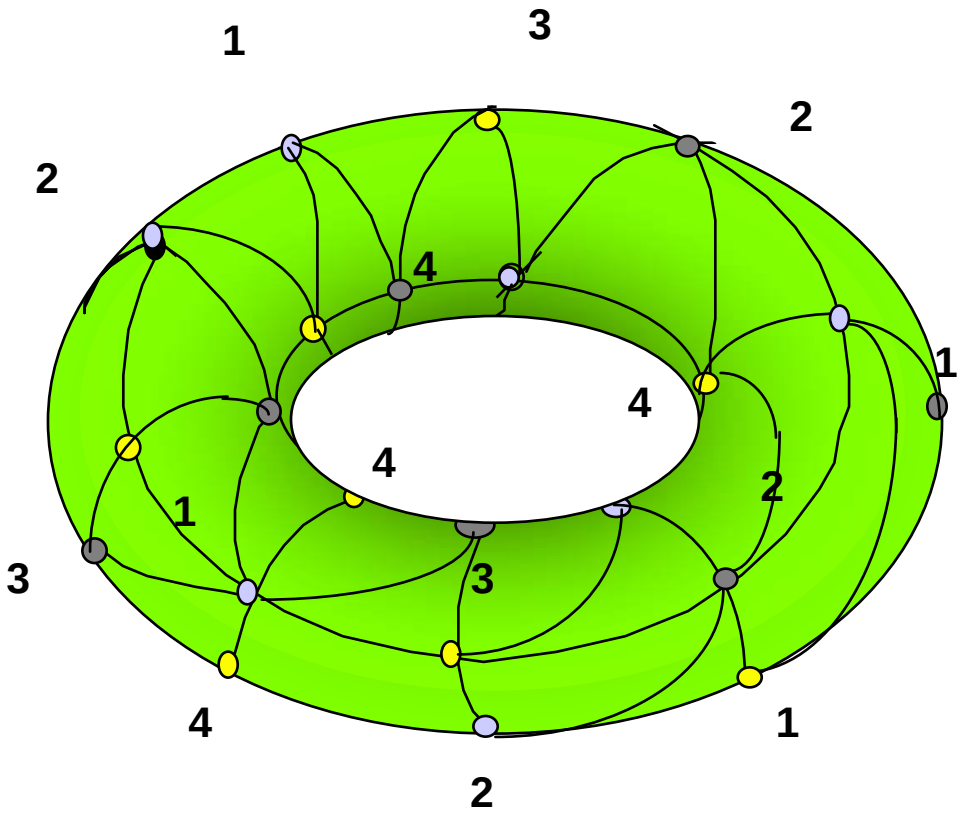

3 Planes, 4 Slots

input complex

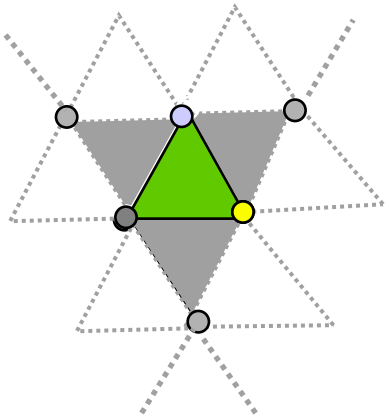

output complex

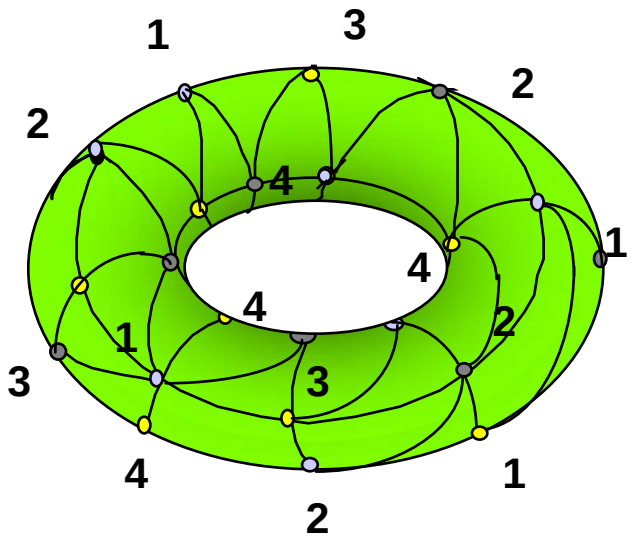

Shared Memory

- Asynchronous
  - arbitrary delays
  - e.g., interrupts, page faults, etc.
- wait-free
  - Processes can fail or be slow
- Communication by reading and writing shared memory

Asynchronous Computability Theorem

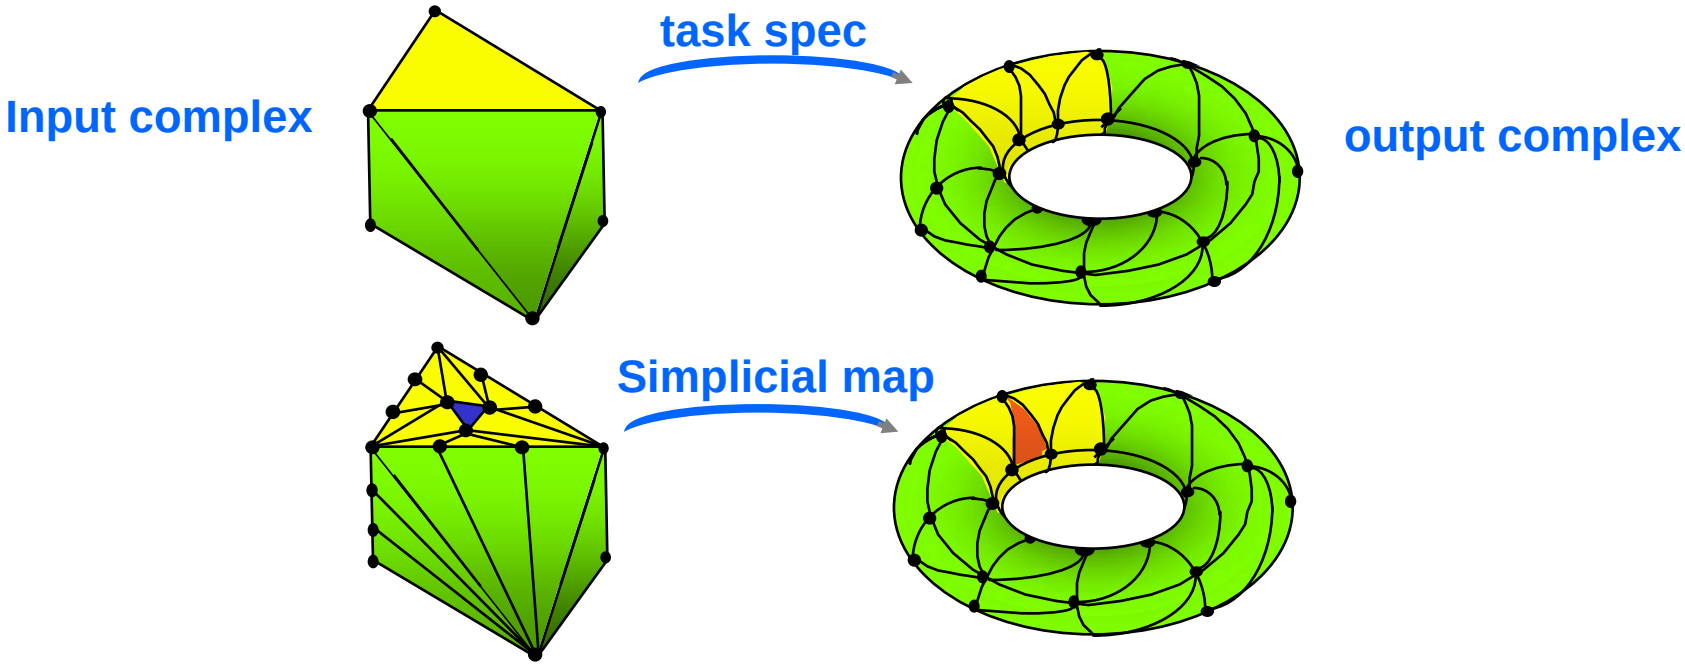

A task has a ***wait-free*** protocol if and only if one can ***chromatically subdivide*** its input complex so that there exists a ***color preserving simplicial map*** to its output complex that refines the task spec.

Protocol Complex

- Each protocol defines a complex
  - vertex: sequence of messages received
  - simplex: compatible set of vertexes
- Induces subdivision of input complex

One-Round Protocol Complex

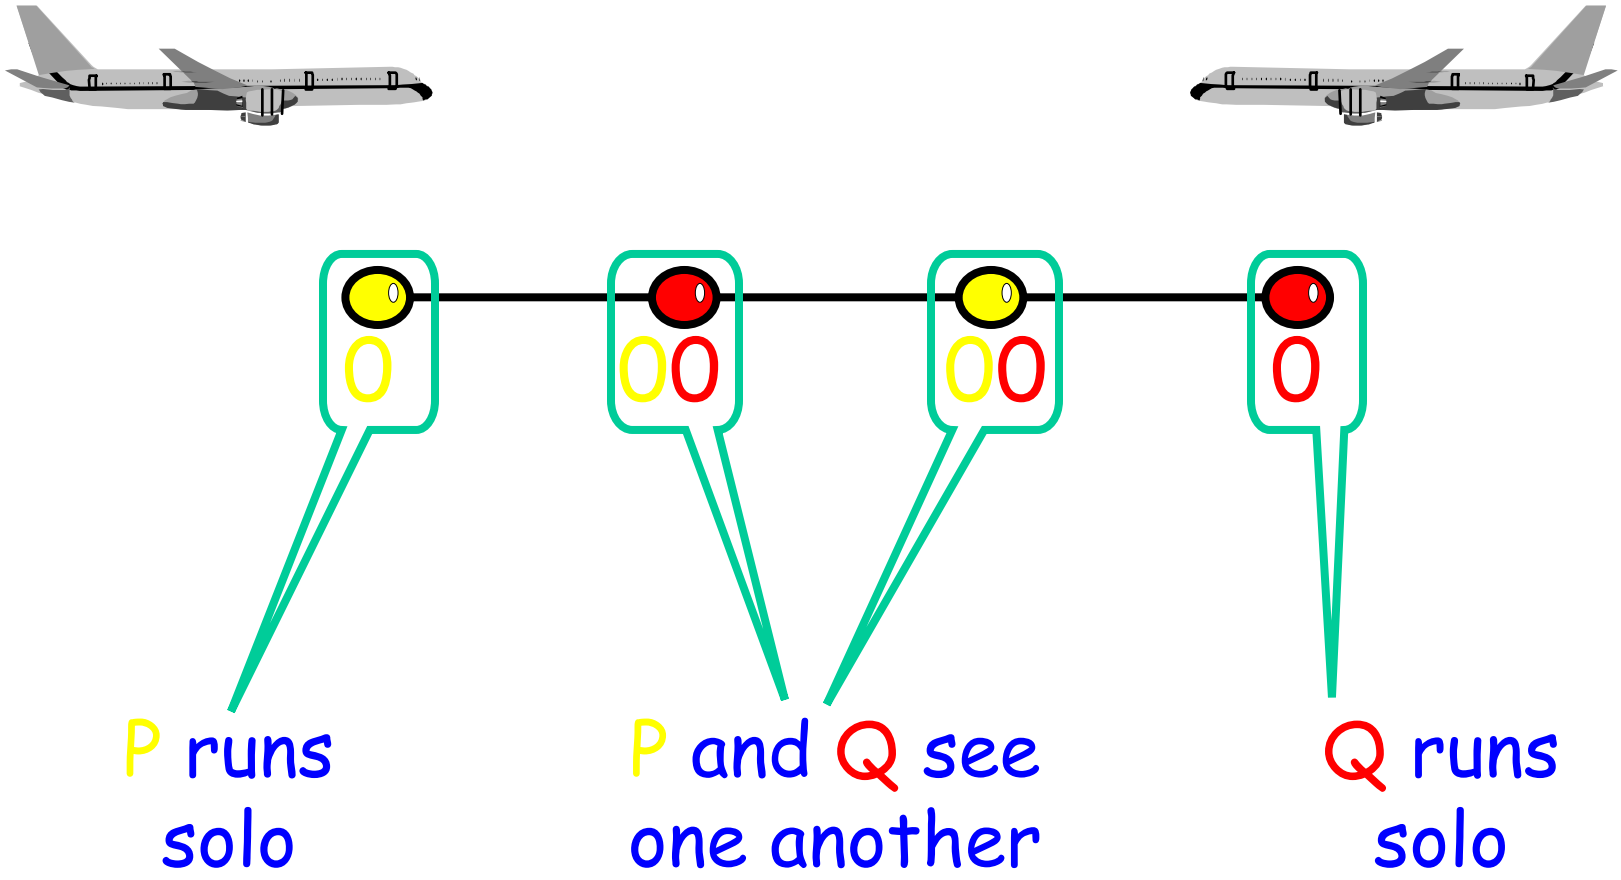

One-Round Protocol Complex

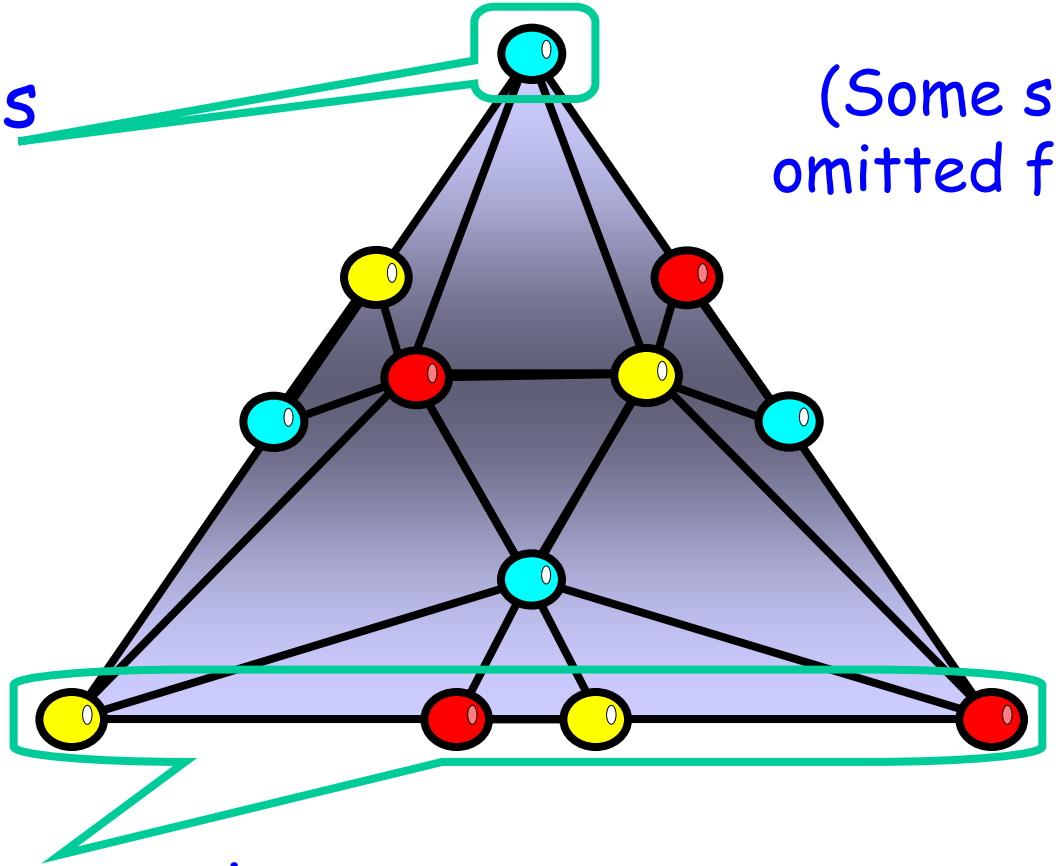

Summary

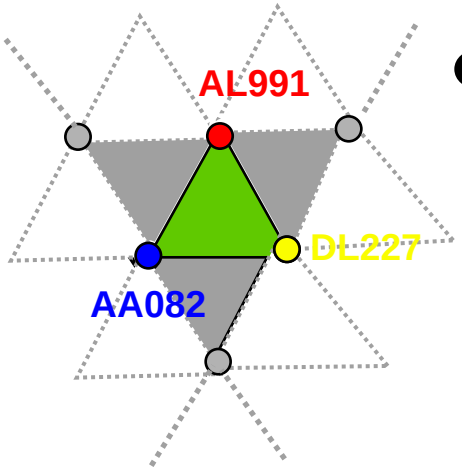

Input complex

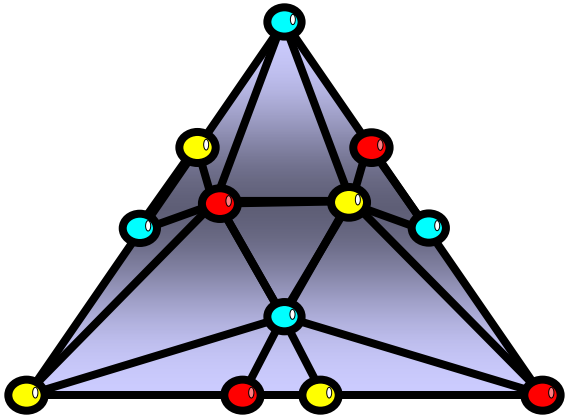

Protocol complex

δ

Task def

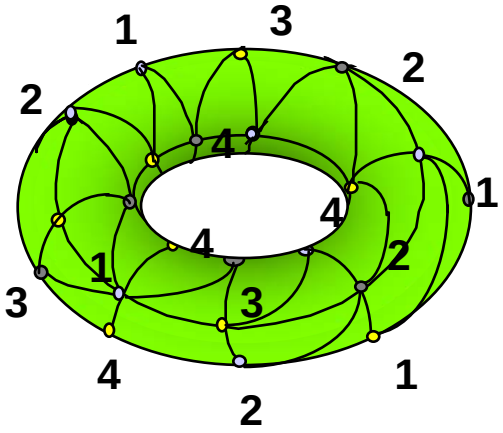

Output complex

2 Planes 3 Slots

```
If view = (my_id, ⊥)
  decide 1
  else if my_id < his_id
       decide 2
       else decide 3
```

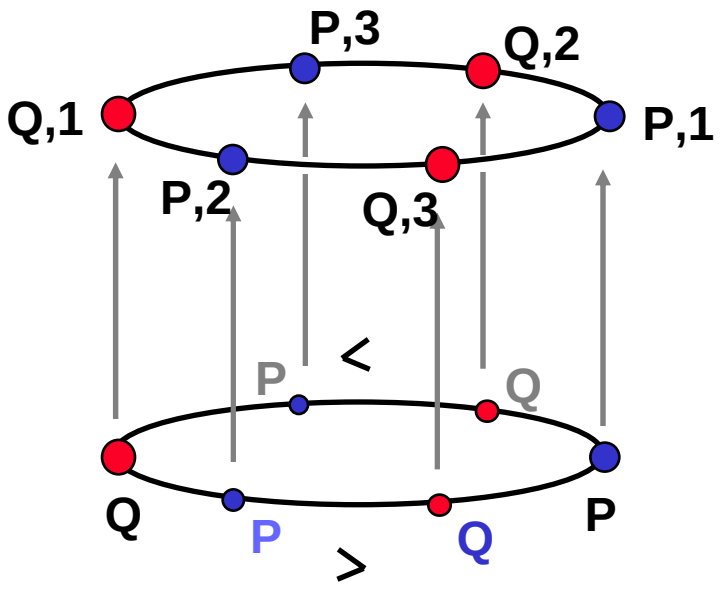

simplicial map

comparison based algorithm: only compares plane ids

3 Planes 4 Slots

- Impossibility by reduction
- Assume a protocol for 3 airplanes, 4 slots
- Choose
  - 0 if your name is even
  - 1 if your name is odd
- Result
  - Not all odd
  - Not all even

Reminder: Cannot Map Boundary Around a Hole

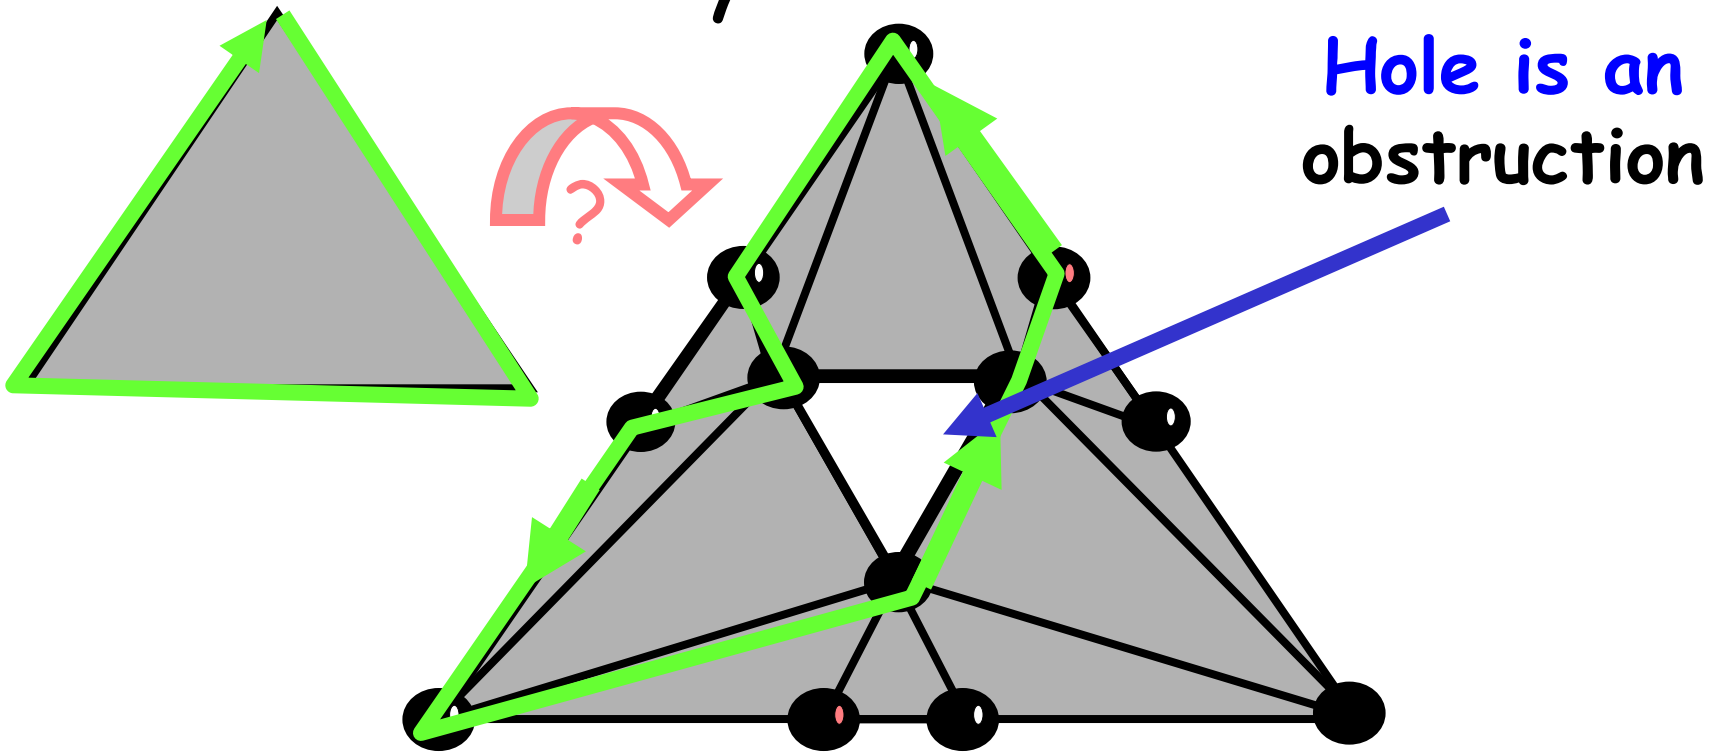

Output Complex (3 processes)

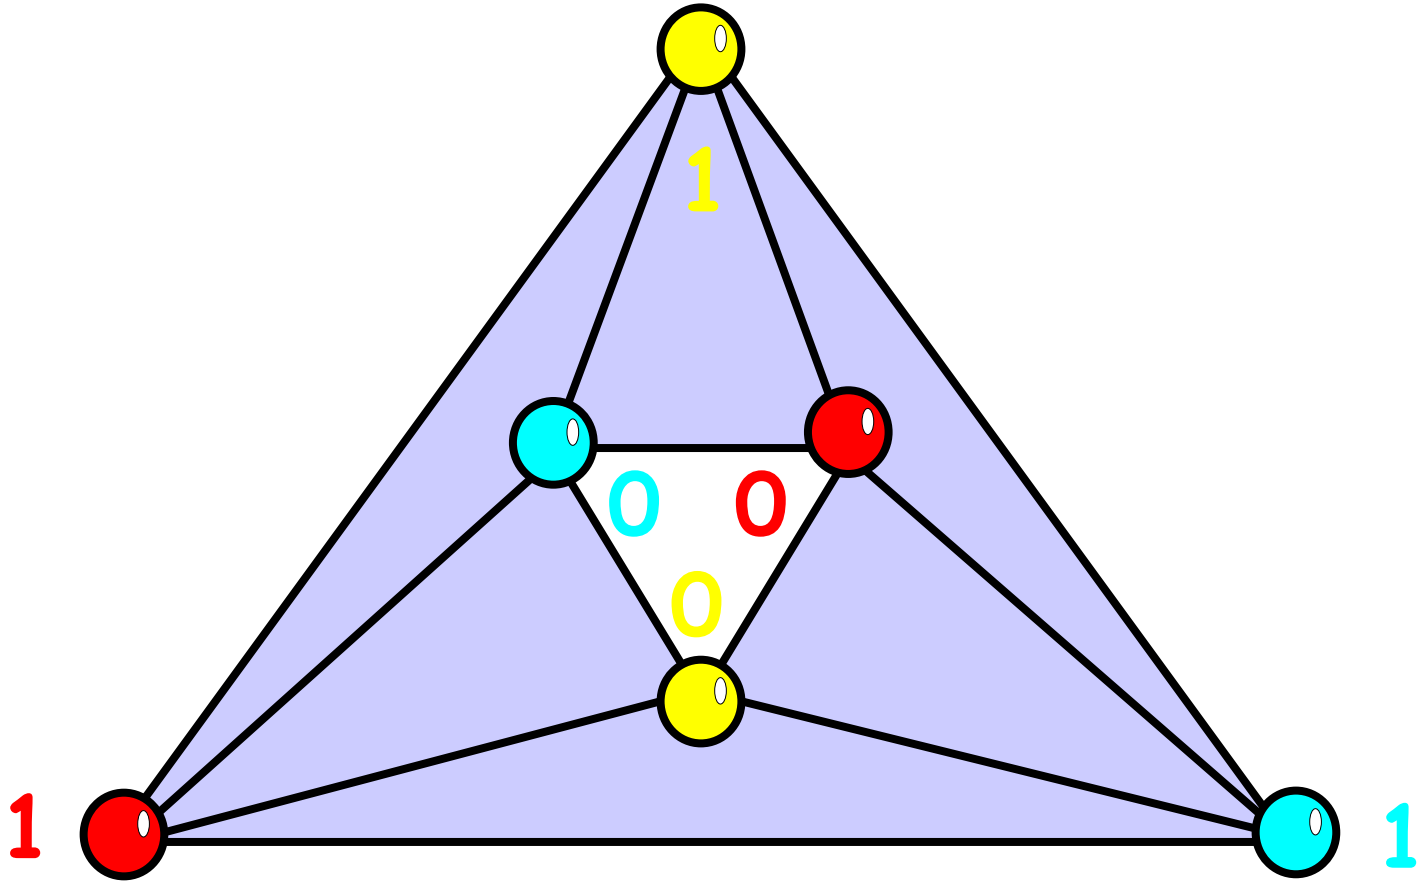

Protocol Complex (schematic)

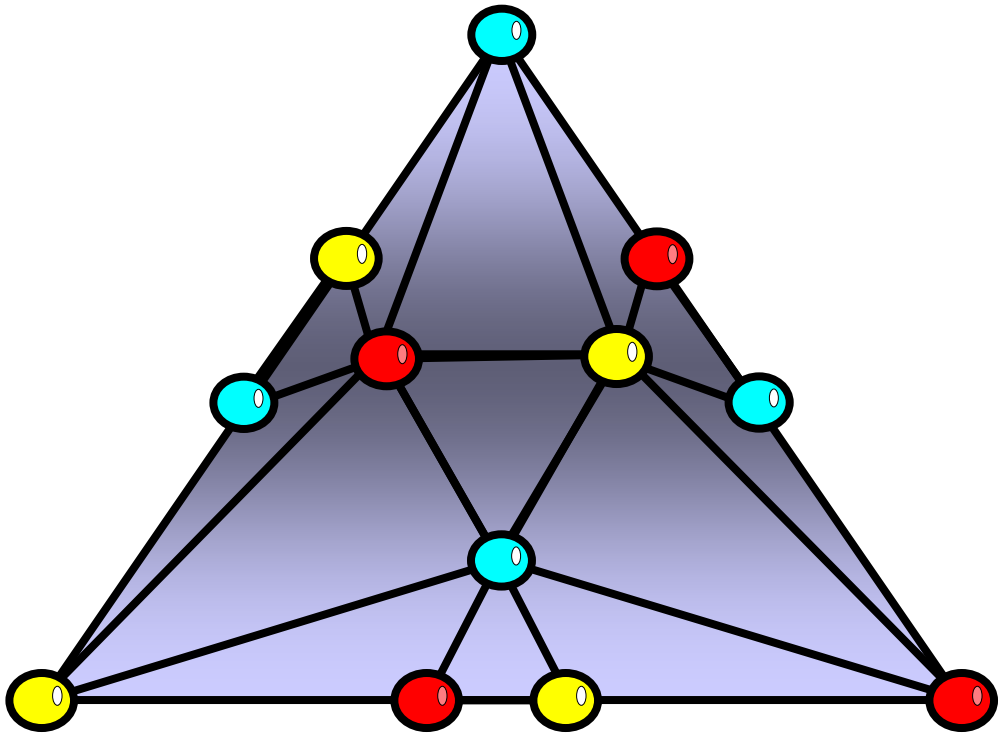

Boundary = 2-Process Executions

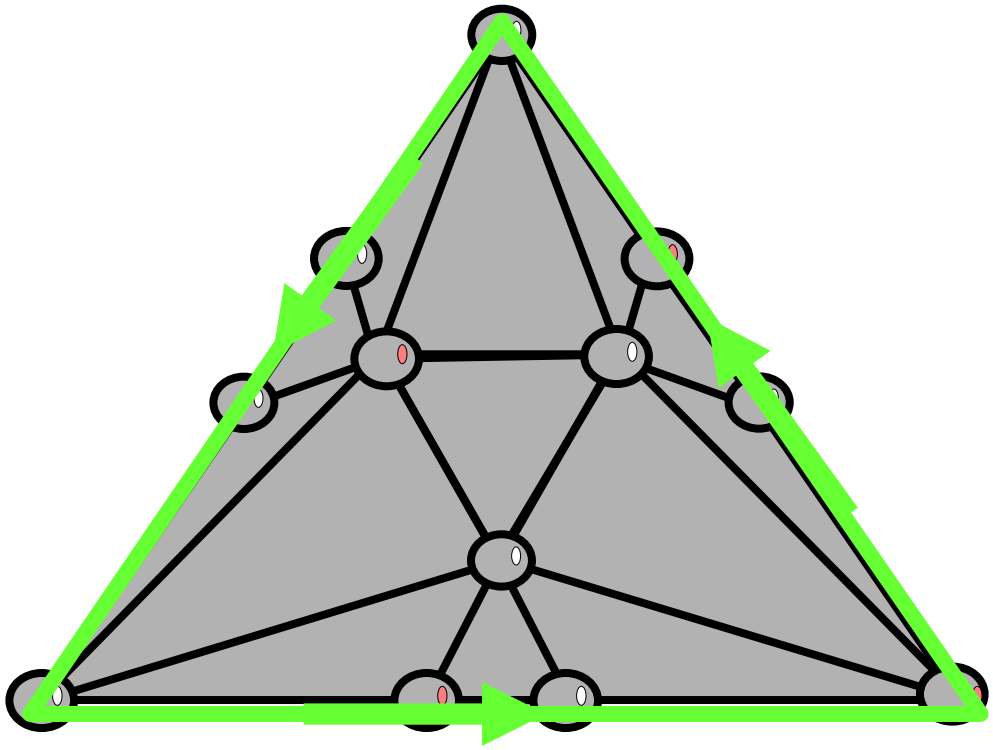

Protocol Complex for One Process Execution

P() decides 1 WLOG

P() decides 1 by symmetry

P() decides 1 by symmetry

2-Process execution might be mapped this way …

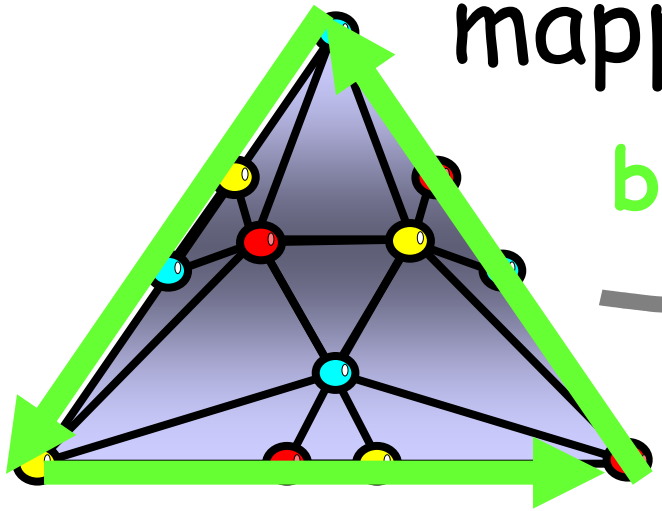

boundary

protocol

Wraps around
-1 times

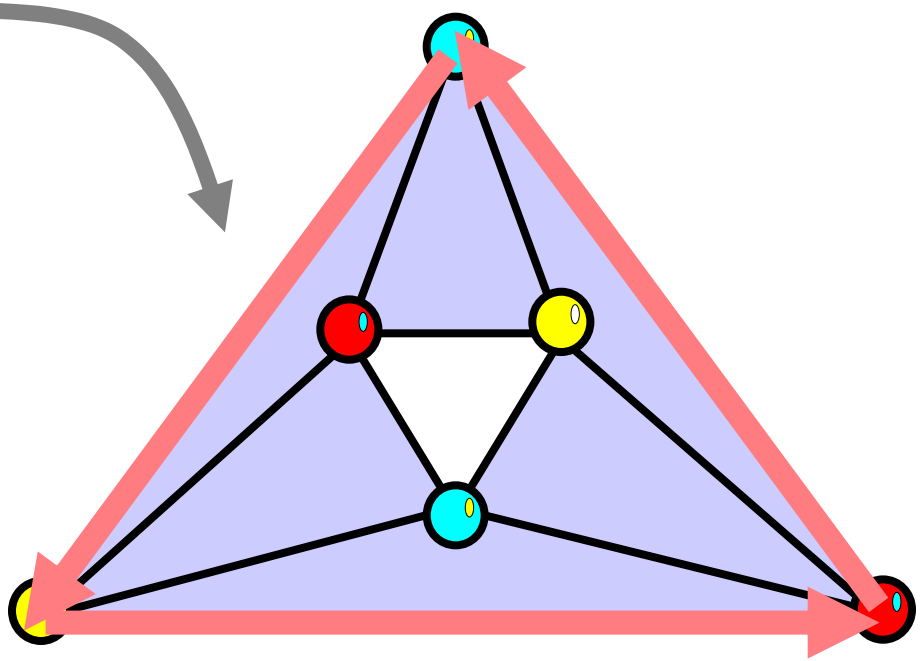

output

2-Process execution might be mapped that way …

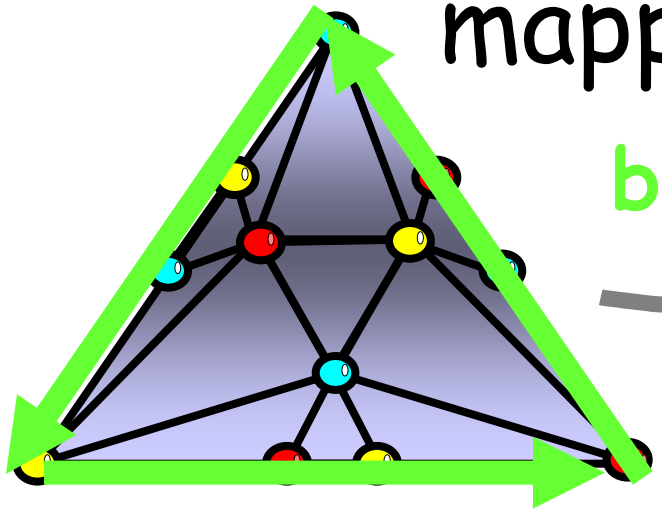

boundary

protocol

Wraps around
+2 times

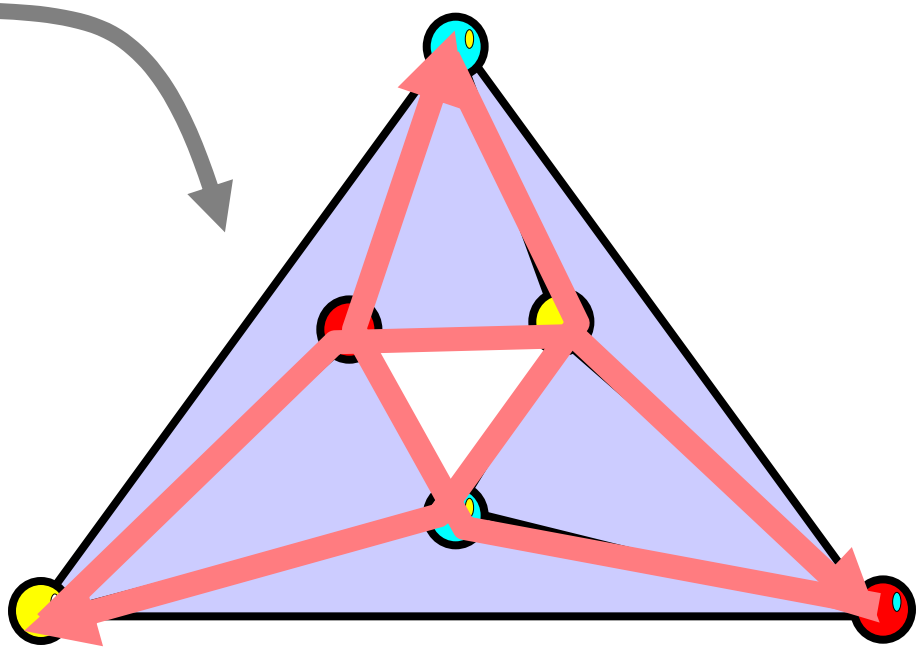

output

In General …

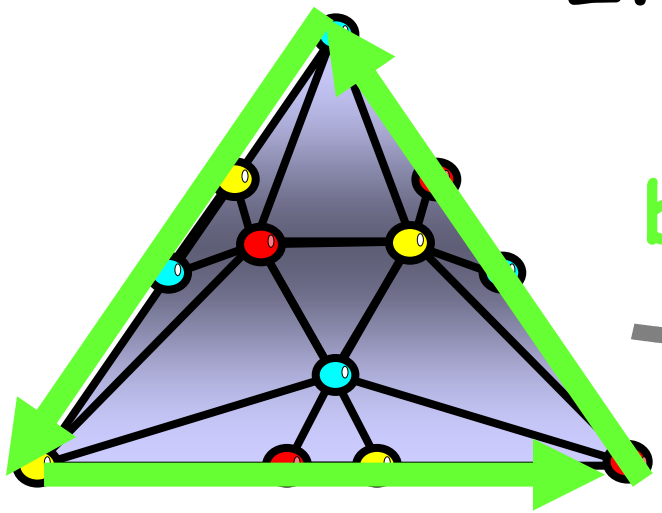

boundary

protocol

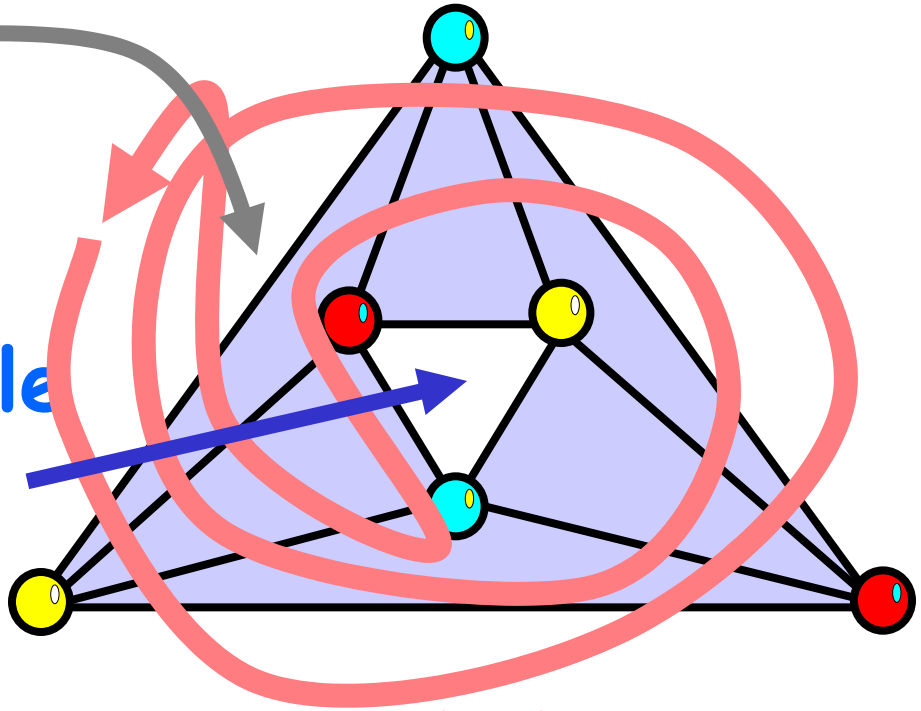

output

Wraps around hole

3k-1 ≠ 0 times

QED!

Other Models of Computation

- Powerful atomic operations
  - Compare&swap, test&set, etc
- Synchronous fail-stop
  - Computation in rounds

3 Planes Need 5 Slots

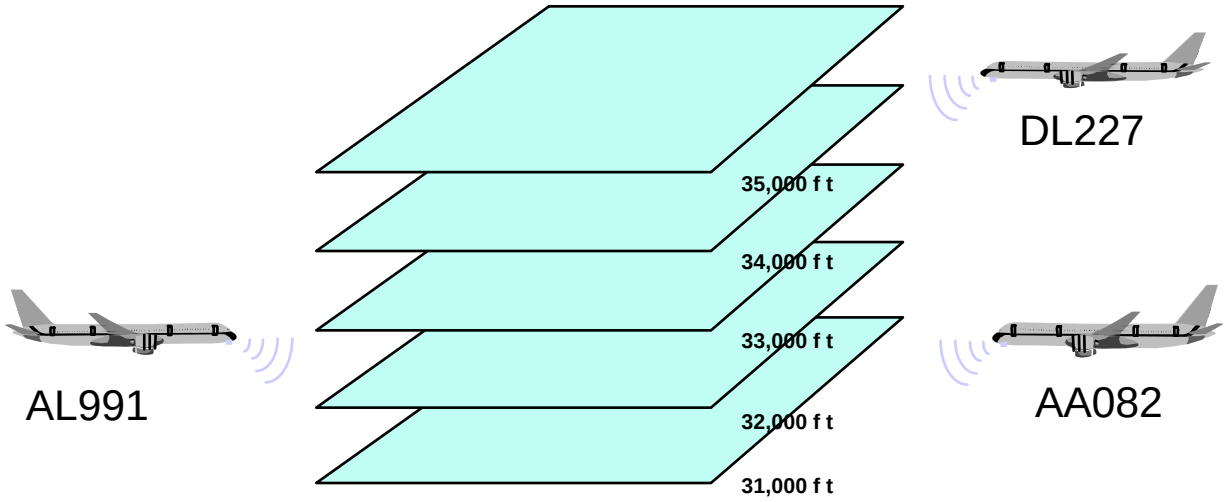

Topological Approach to Asynchrony

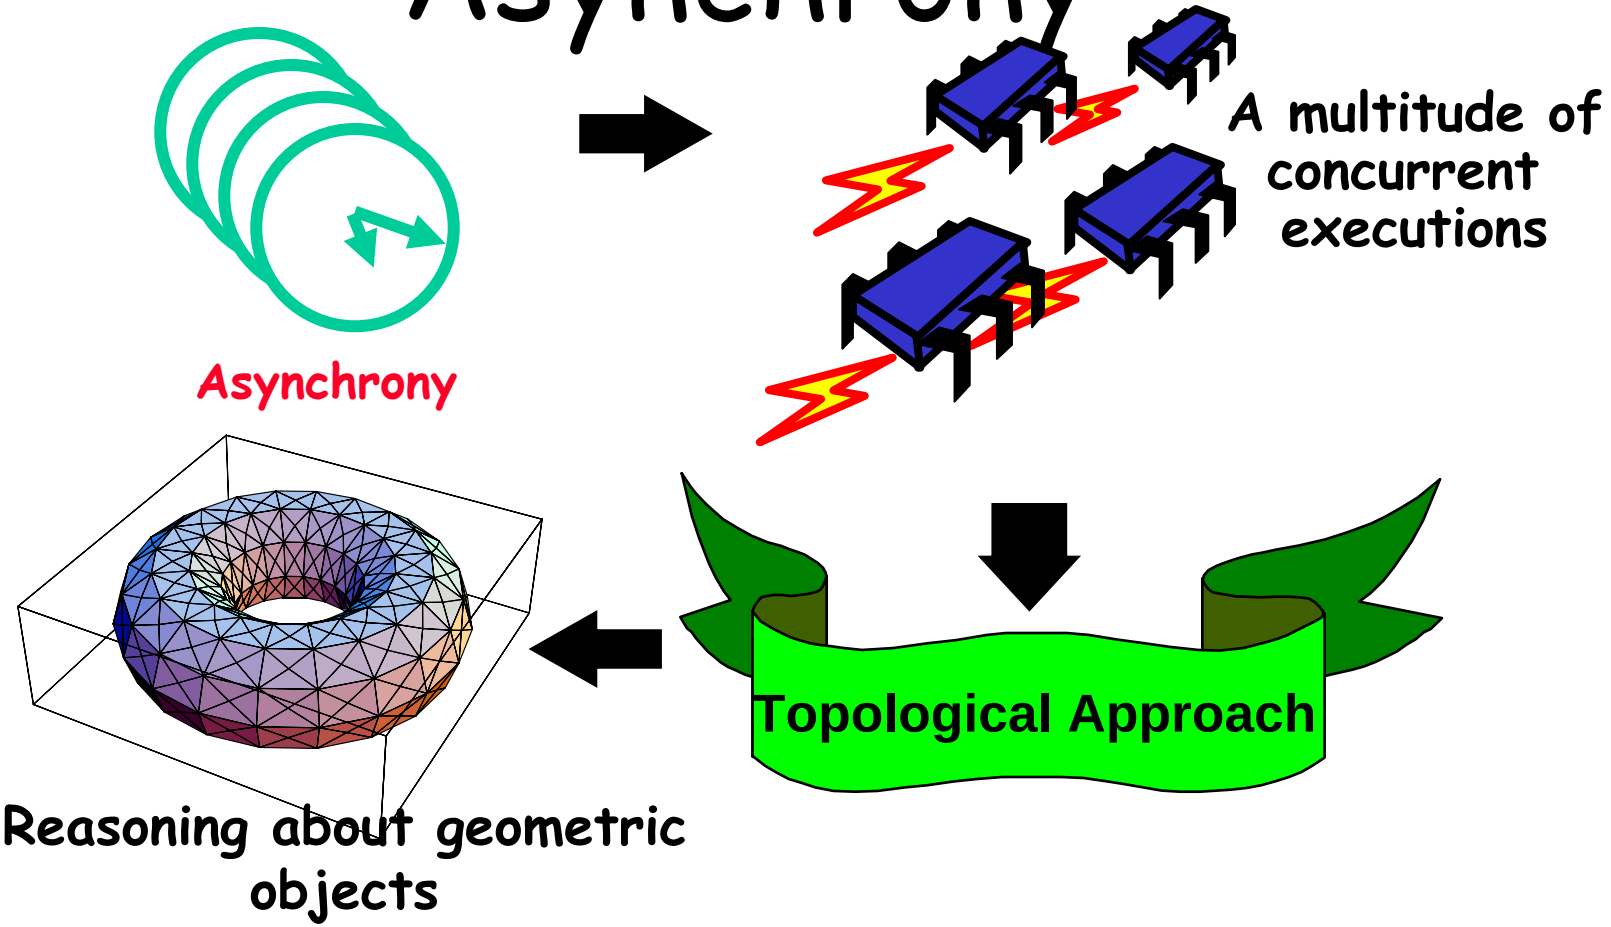

The Glorious Future

- Our work
  - Asynchronous, wait-free, one-shot, RW memory
- Open problems
  - Long-lived computations
  - Other kinds of memory (compare&swap)
  - Randomization
  - Other progress conditions …