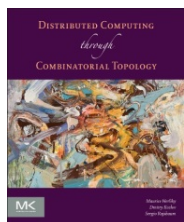


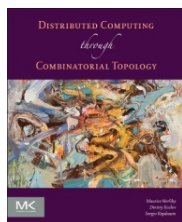
Distributed Computing through Combinatorial Topology

MITRO207, P4, 2019

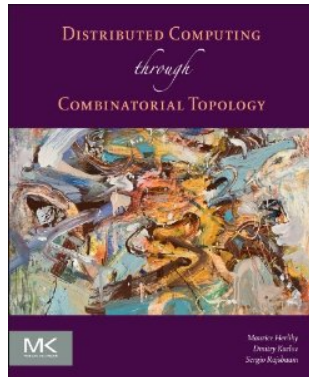


Administrivia

- Language: English? Français sur demande
- Lectures: Wednesday, 8:30-11:45
- Web page: <http://perso.telecom-paristech.fr/~kuznetso/MITRO207-2019/>
- Homeworks
 - Corrected, not graded
 - TD: 29.05.2019
- Office hours:
 - C213-2, appointments by email to petr.kuznetsov@telecom-paristech.fr
- Credit = written exam: June 25, 2019 (3 hours)
- Bonus for homeworks, participation, discussion of exercises, bugs found



Literature

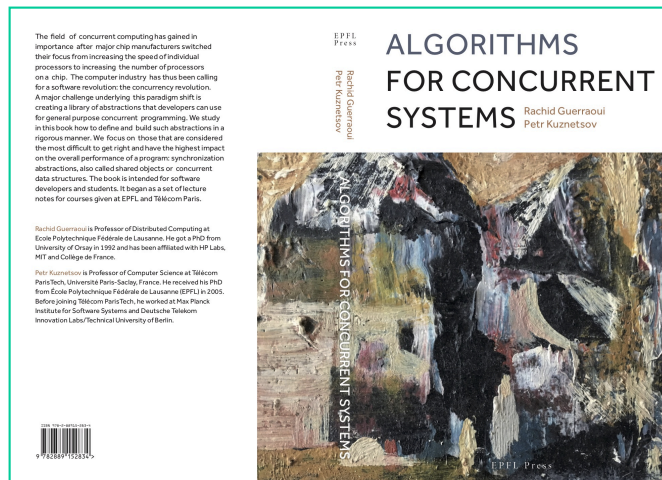


Distributed Computing Through Combinatorial Topology

Maurice Herlihy, Dmitry Kozlov, Sergio Rajsbaum

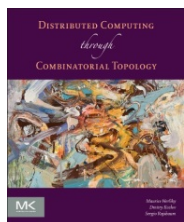
□ Morgan Kaufman, 2013, available online (TPT library)

- Algorithms for Concurrent Systems.
R. Guerraoui, P. Kuznetsov, 2018

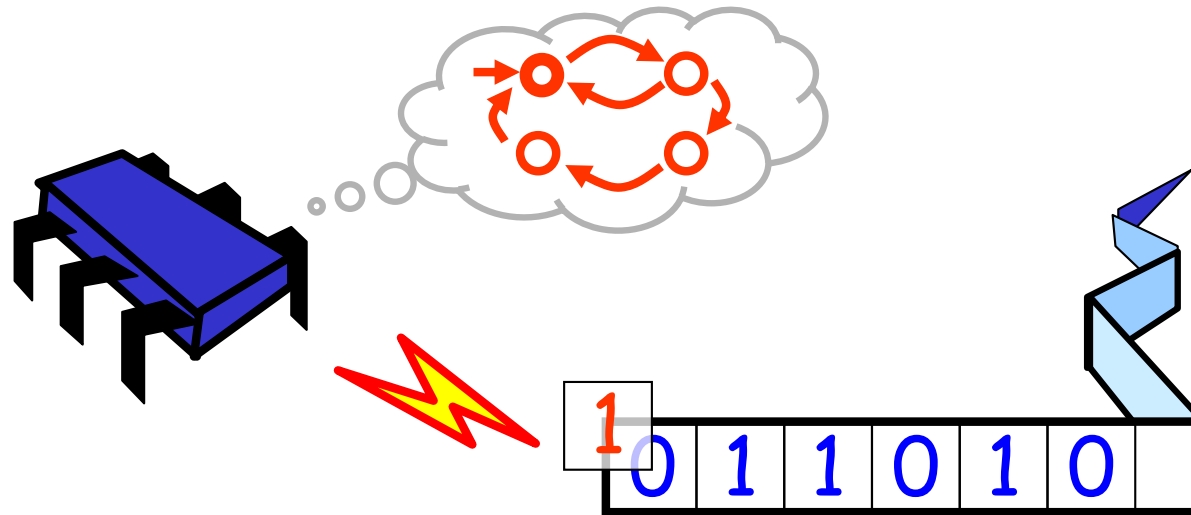


(Preliminary) road map

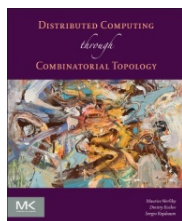
- The matter and the method of distributed computing
- Basics of combinatorial topology
- Colorless tasks
- Simulations and reductions
- Generic tasks and manifold computations
- Renaming and oriented manifolds



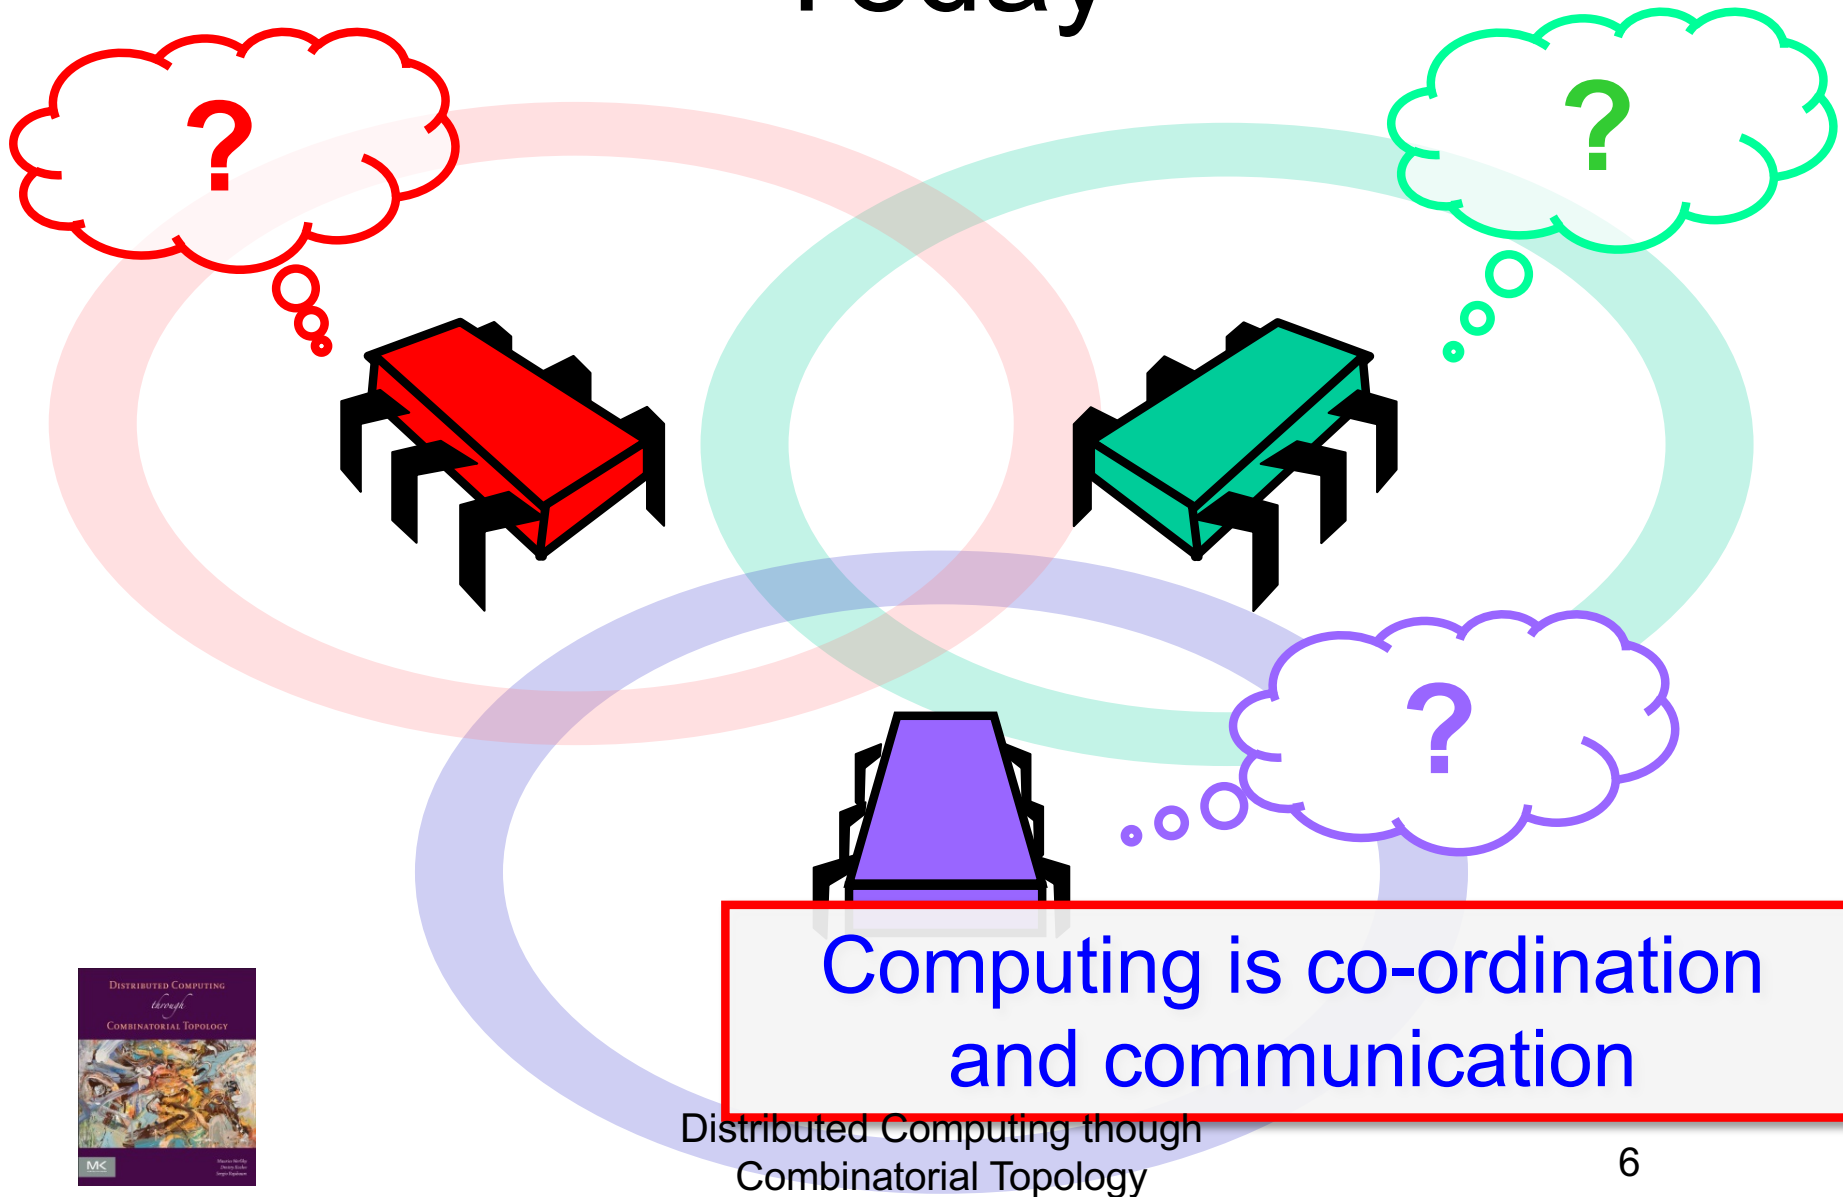
In the Beginning ...

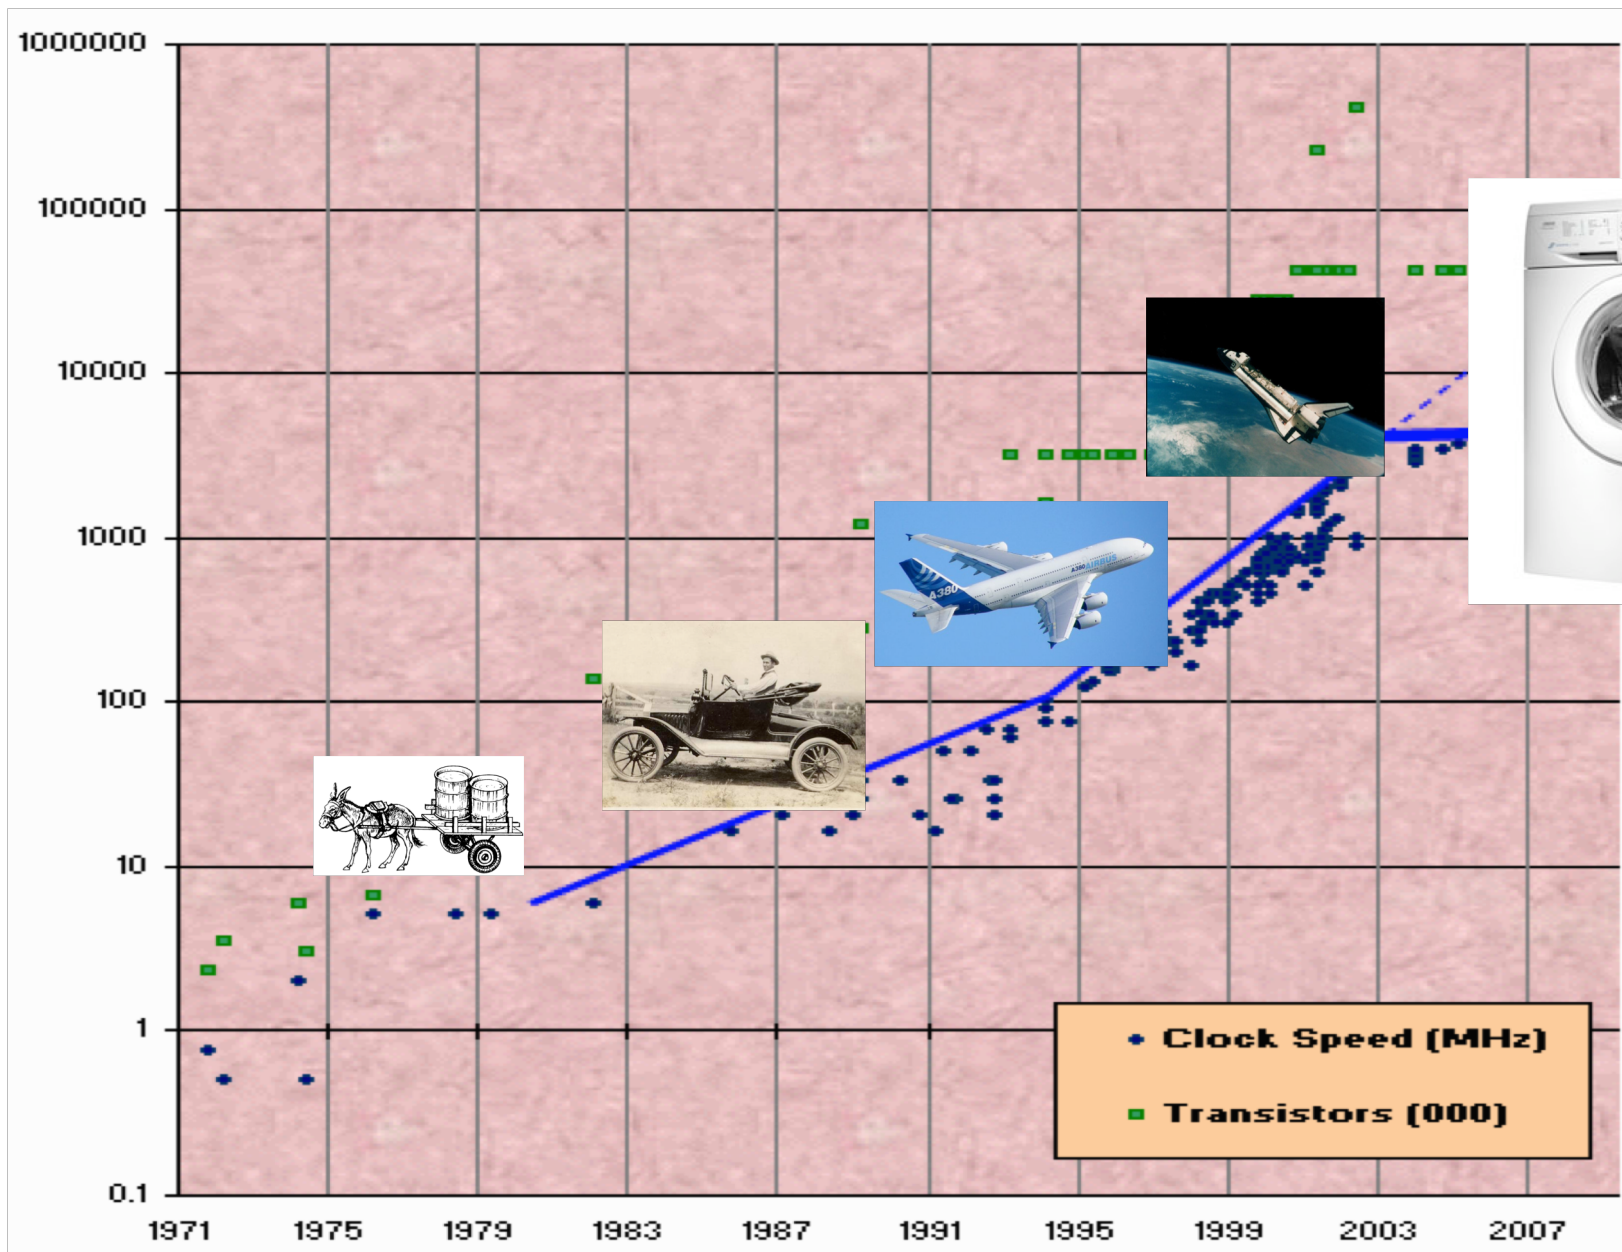


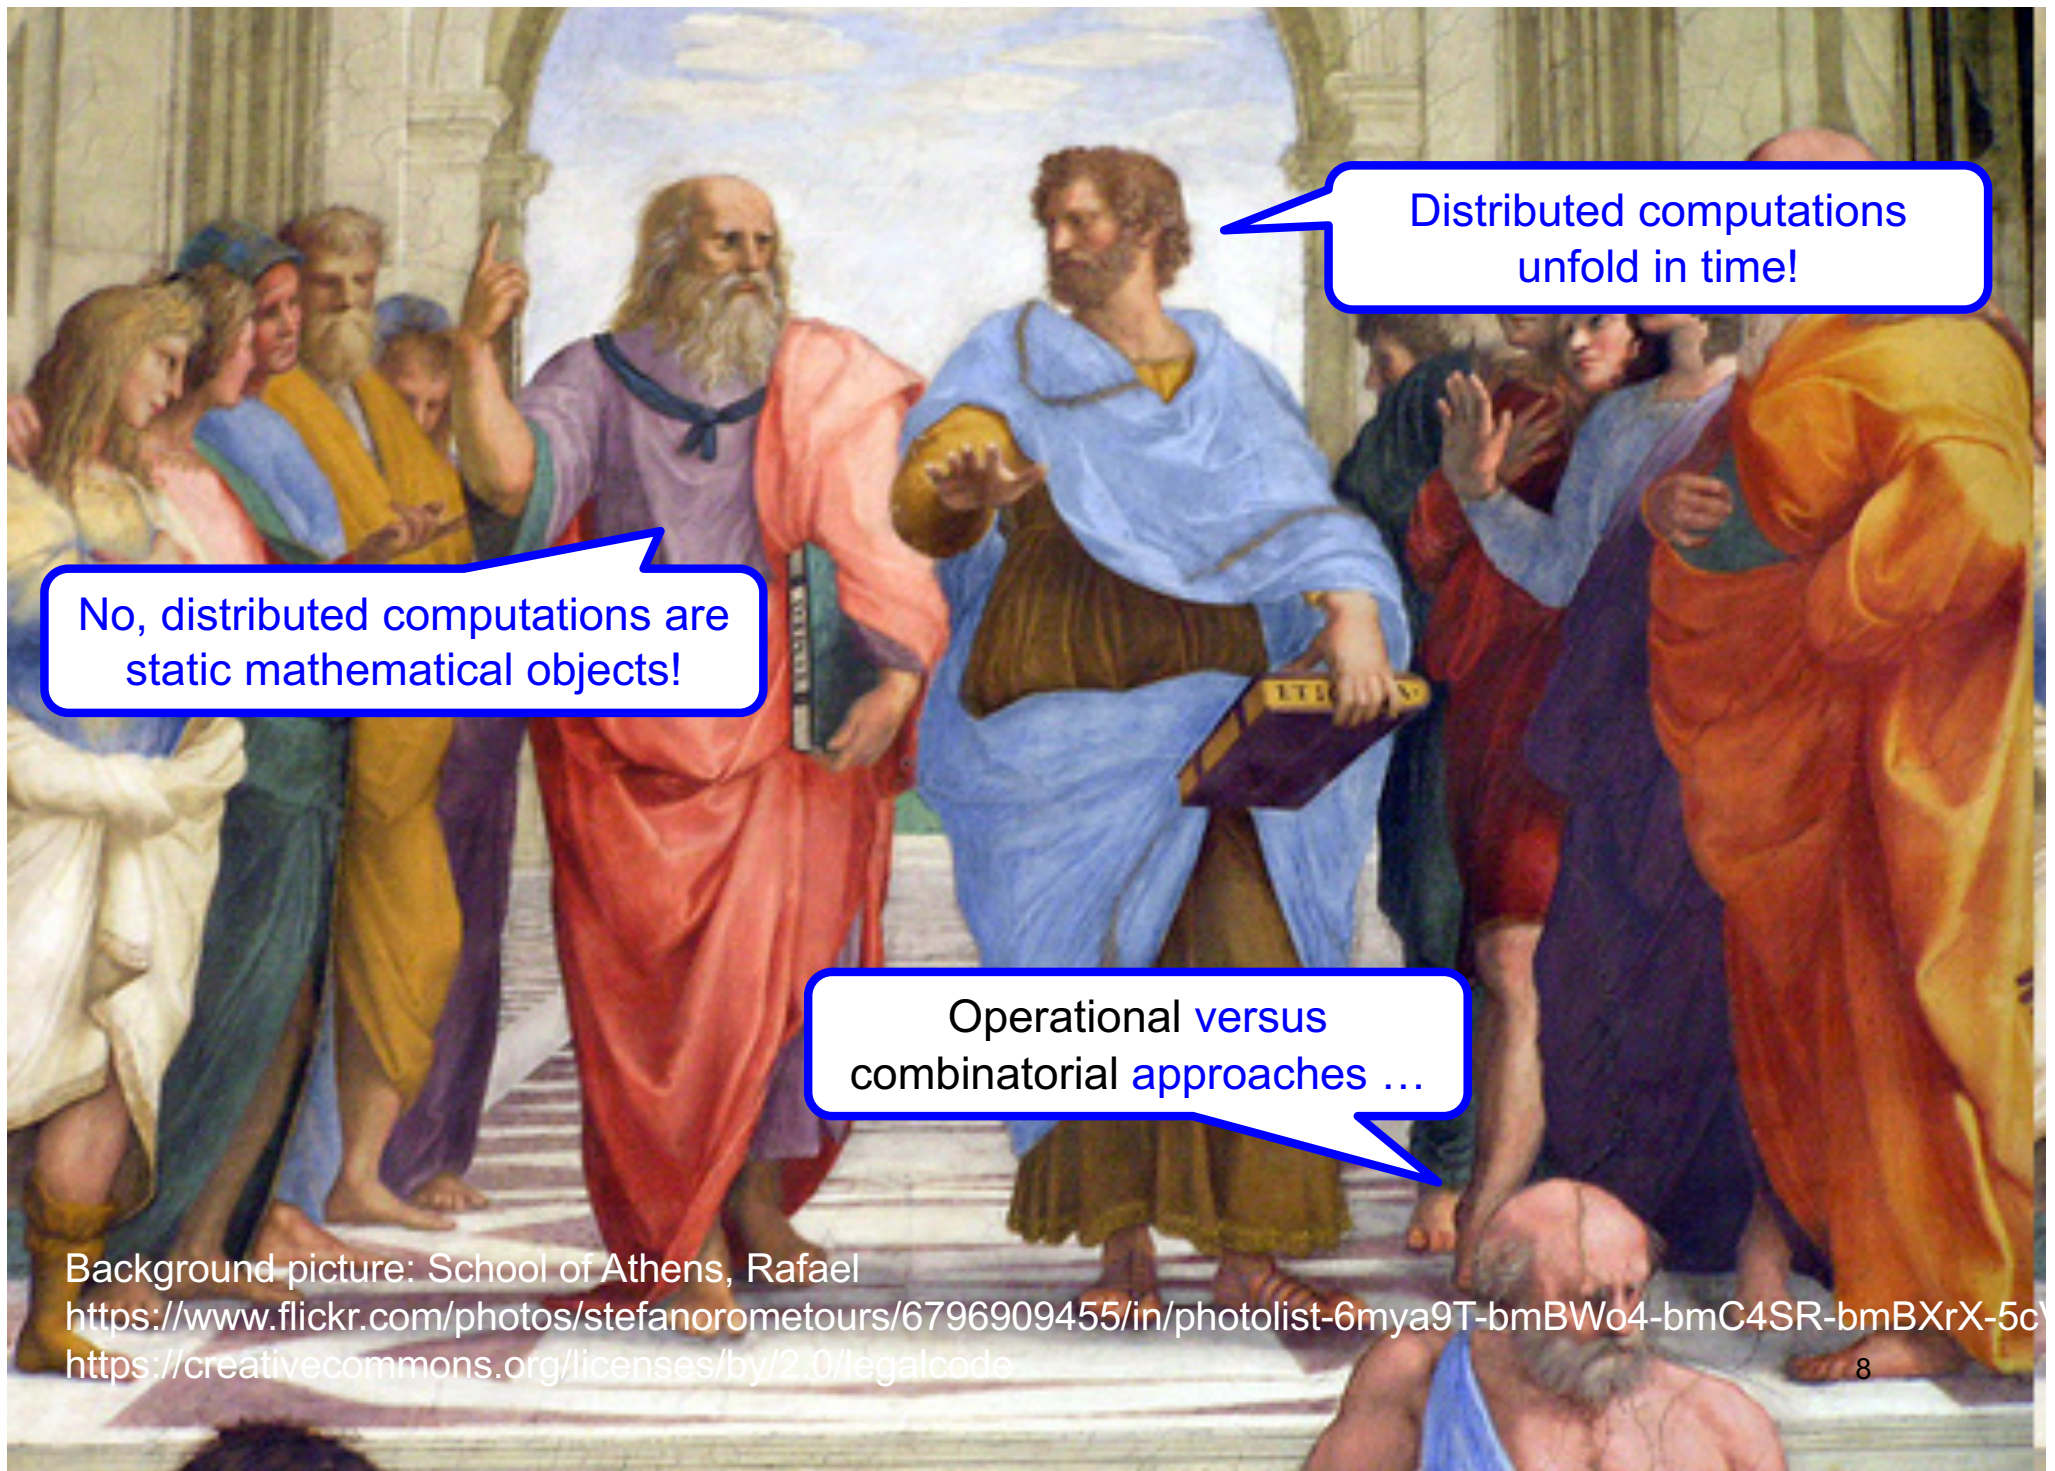
a computer was just a Turing machine ...



Today







No, distributed computations are static mathematical objects!

Distributed computations unfold in time!

Operational versus combinatorial approaches ...

Background picture: School of Athens, Rafael

<https://www.flickr.com/photos/stefanorometours/6796909455/in/photolist-6mya9T-bmBW04-bmC4SR-bmBXrX-5c>

<https://creativecommons.org/licenses/by/2.0/legalcode>

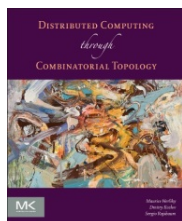
Road Map

Distributed Computing

Two Classic Distributed Problems

The Muddy Children

Coordinated Attack



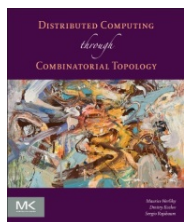
Road Map

Distributed Computing

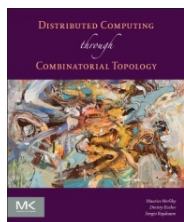
Two Classic Distributed Problems

The Muddy Children

Coordinated Attack

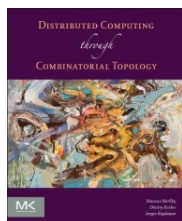


There are Many Models



Distributed Computing through
Combinatorial Topology

There are Many Models



Distributed Computing through
Combinatorial Topology

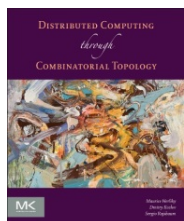
There are Many Models



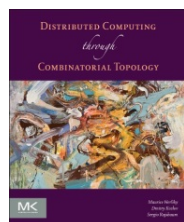
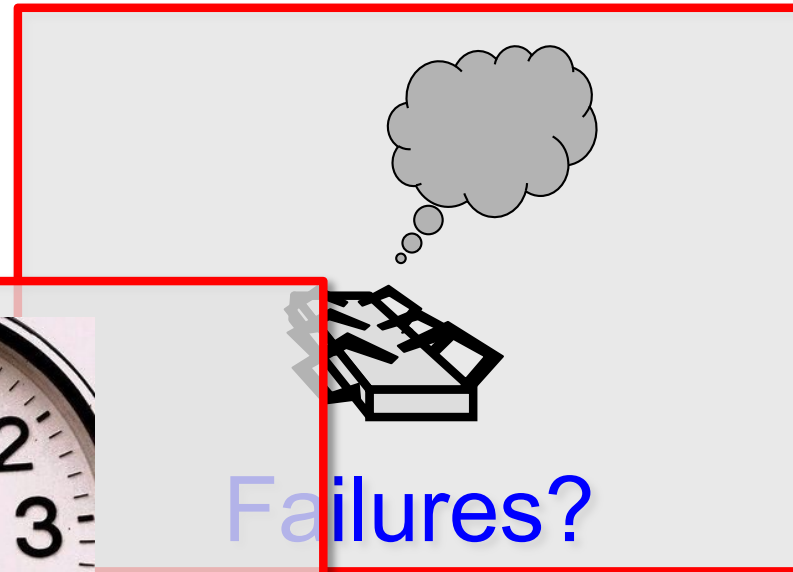
Communication?



Failures?



There are Many Models



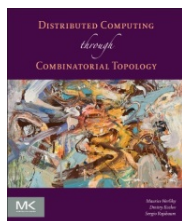
Distributed Computing through
Combinatorial Topology

Communication

<http://commons.wikimedia.org/wiki/File:Pennyblack-pd.jpg>



Message-Passing



Distributed Computing through
Combinatorial Topology

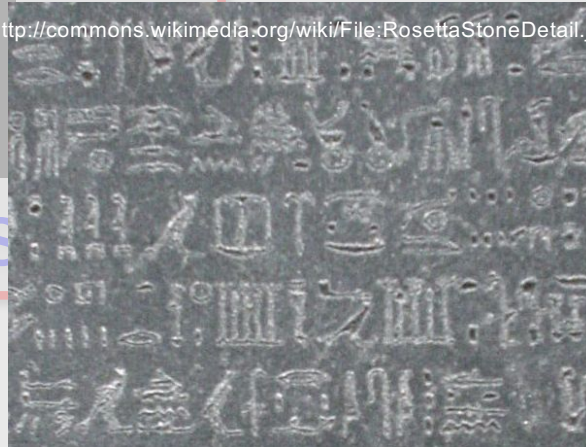
Communication

<http://commons.wikimedia.org/wiki/File:Pennyblack-pd.jpg>

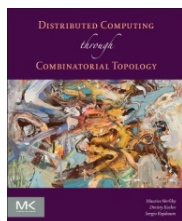


Message-Pas

<http://commons.wikimedia.org/wiki/File:RosettaStoneDetail.jpg>



Read-Write Memory



Distributed Computing through
Combinatorial Topology

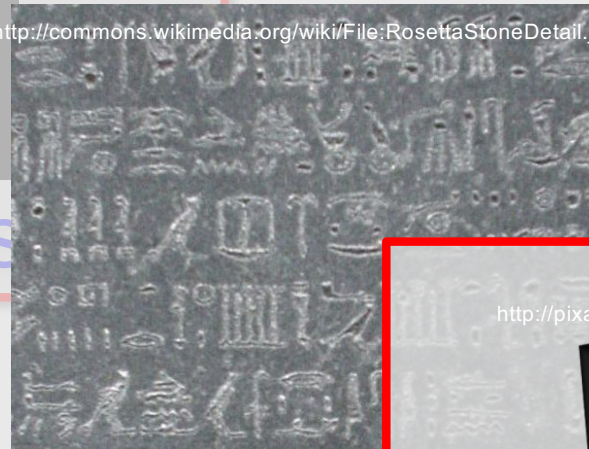
Communication

<http://commons.wikimedia.org/wiki/File:Pennyblack-pd.jpg>



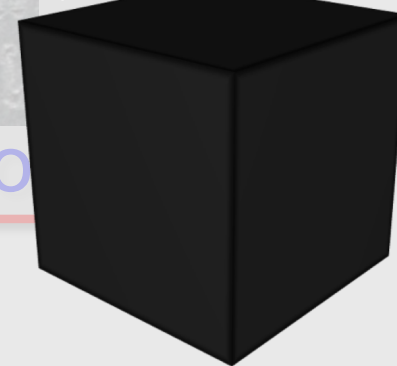
Message-Pas

<http://commons.wikimedia.org/wiki/File:RosettaStoneDetail.jpg>

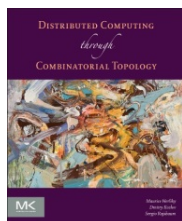


Read-Write Memo

<http://pixabay.com/en/cube-black-black-box-3d-250082/>

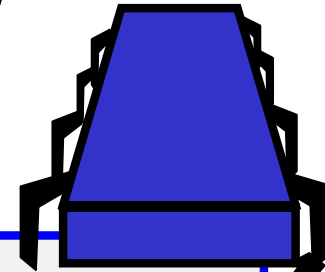


Black-Box Memory



Distributed Computing through
Combinatorial Topology

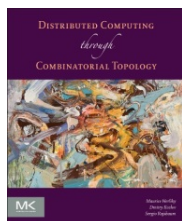
Message-Passing



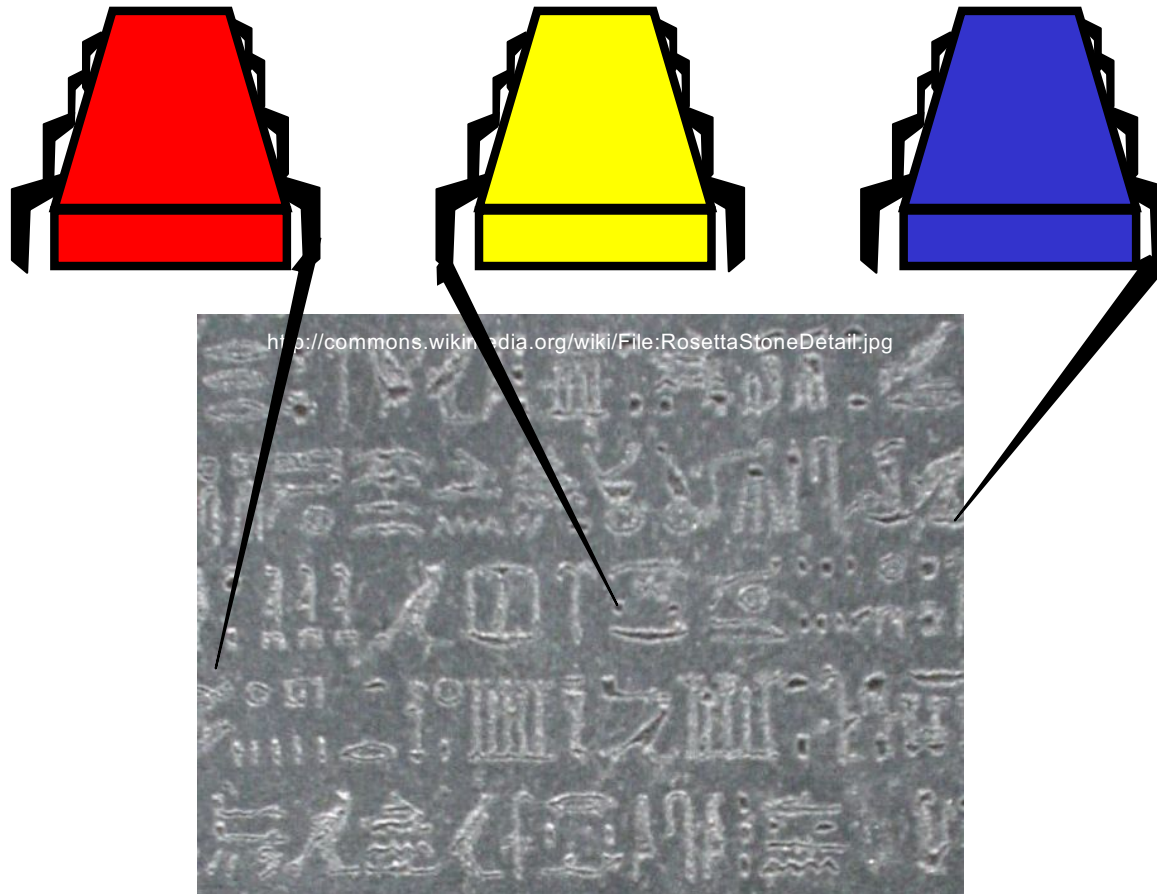
Prof. James Moriarty
Brown University
Providence RI 02912



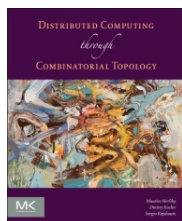
Mr. S. Holmes
221B Baker Street
London NW1 6XE



Read-Write Memory



<http://www.alpa.ch/en/news/2011/dead-sea-scrolls>



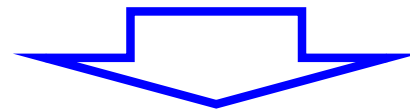
Distributed Computing through
Combinatorial Topology

Read-Write Models

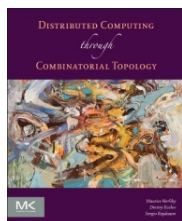
write & read individual locations



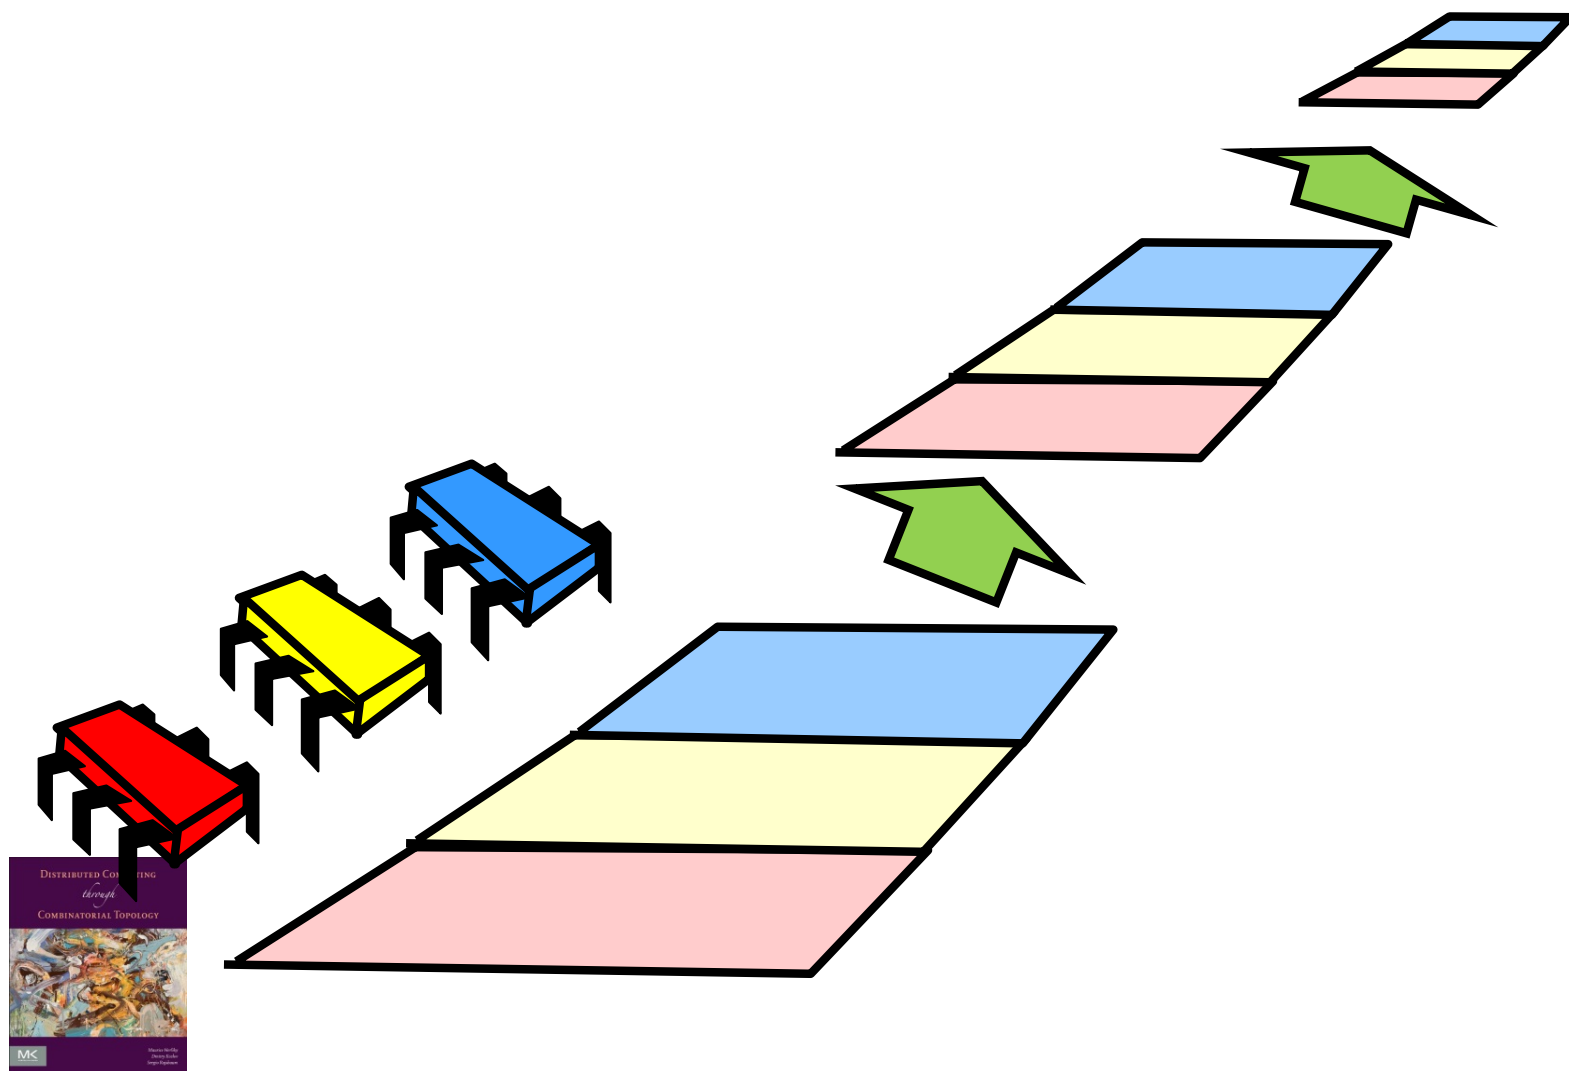
write & take memory snapshot



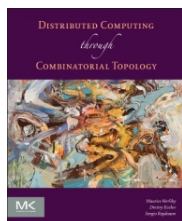
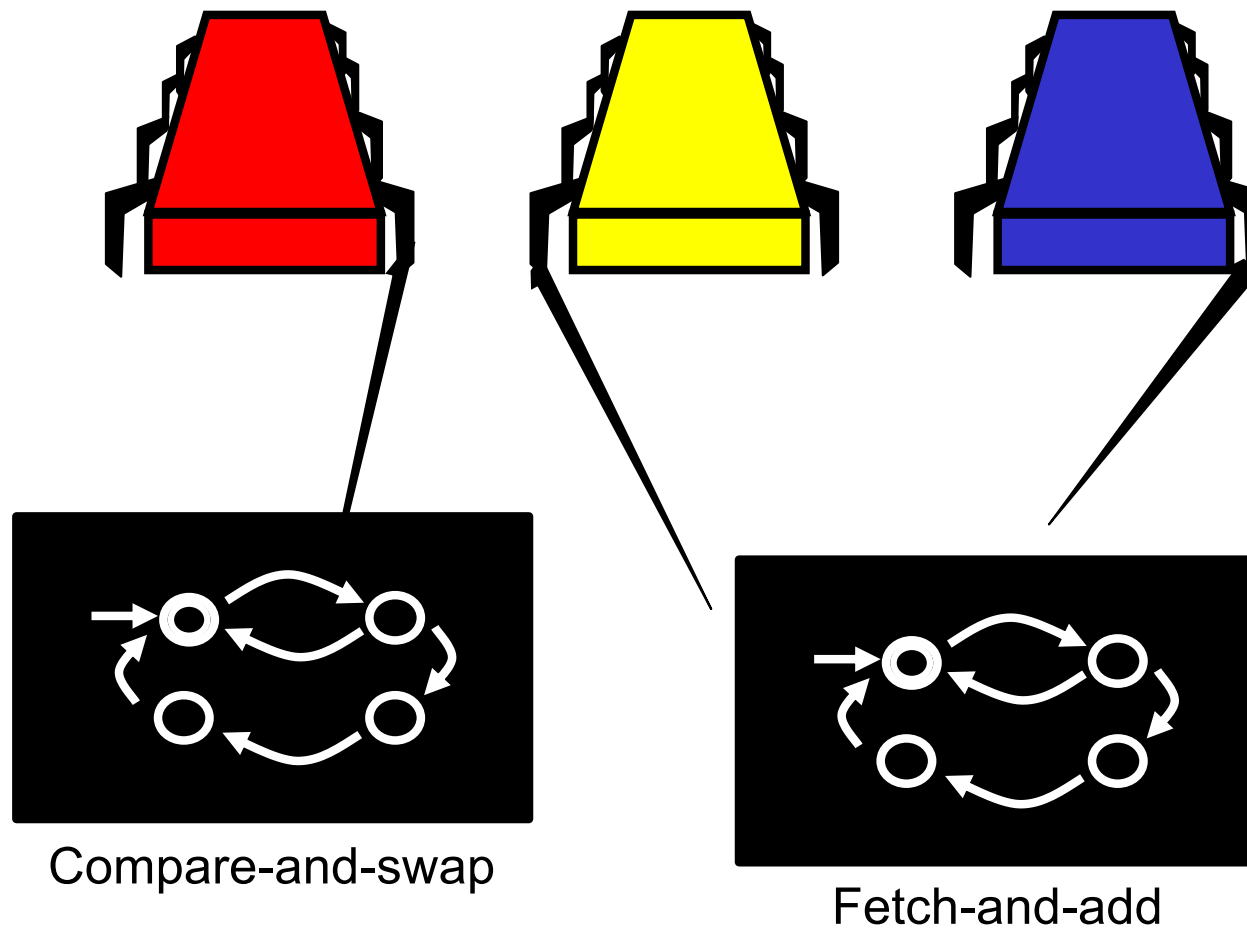
Group writes together then
takes snapshots together



Layered Read-Write Memory



Black Box Memory

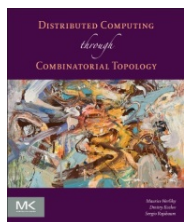
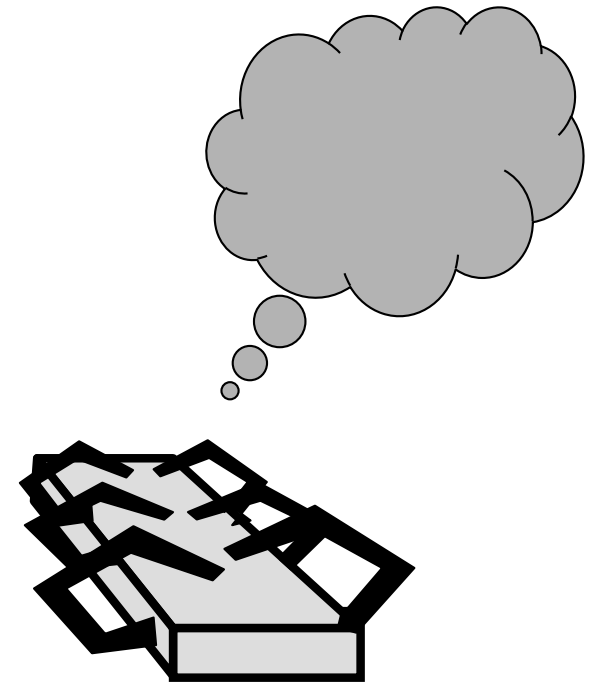


Failures

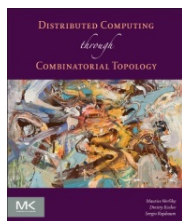
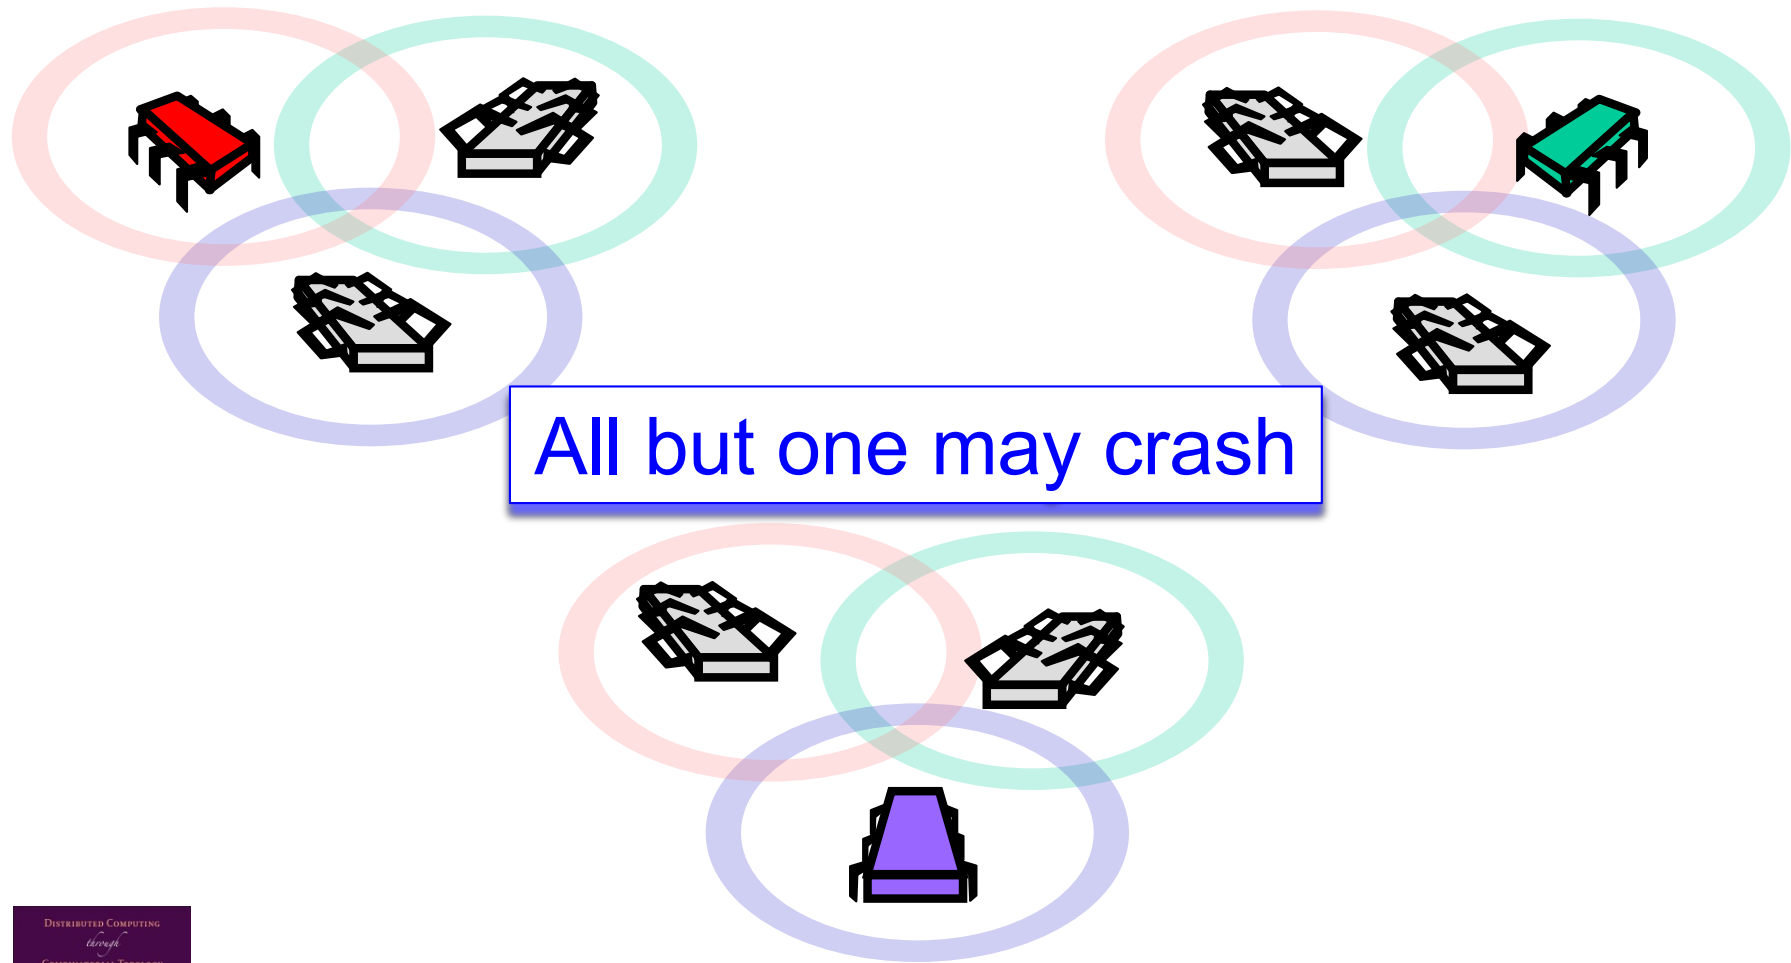
Crash failures: processes halt

How many?

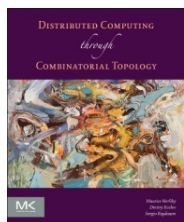
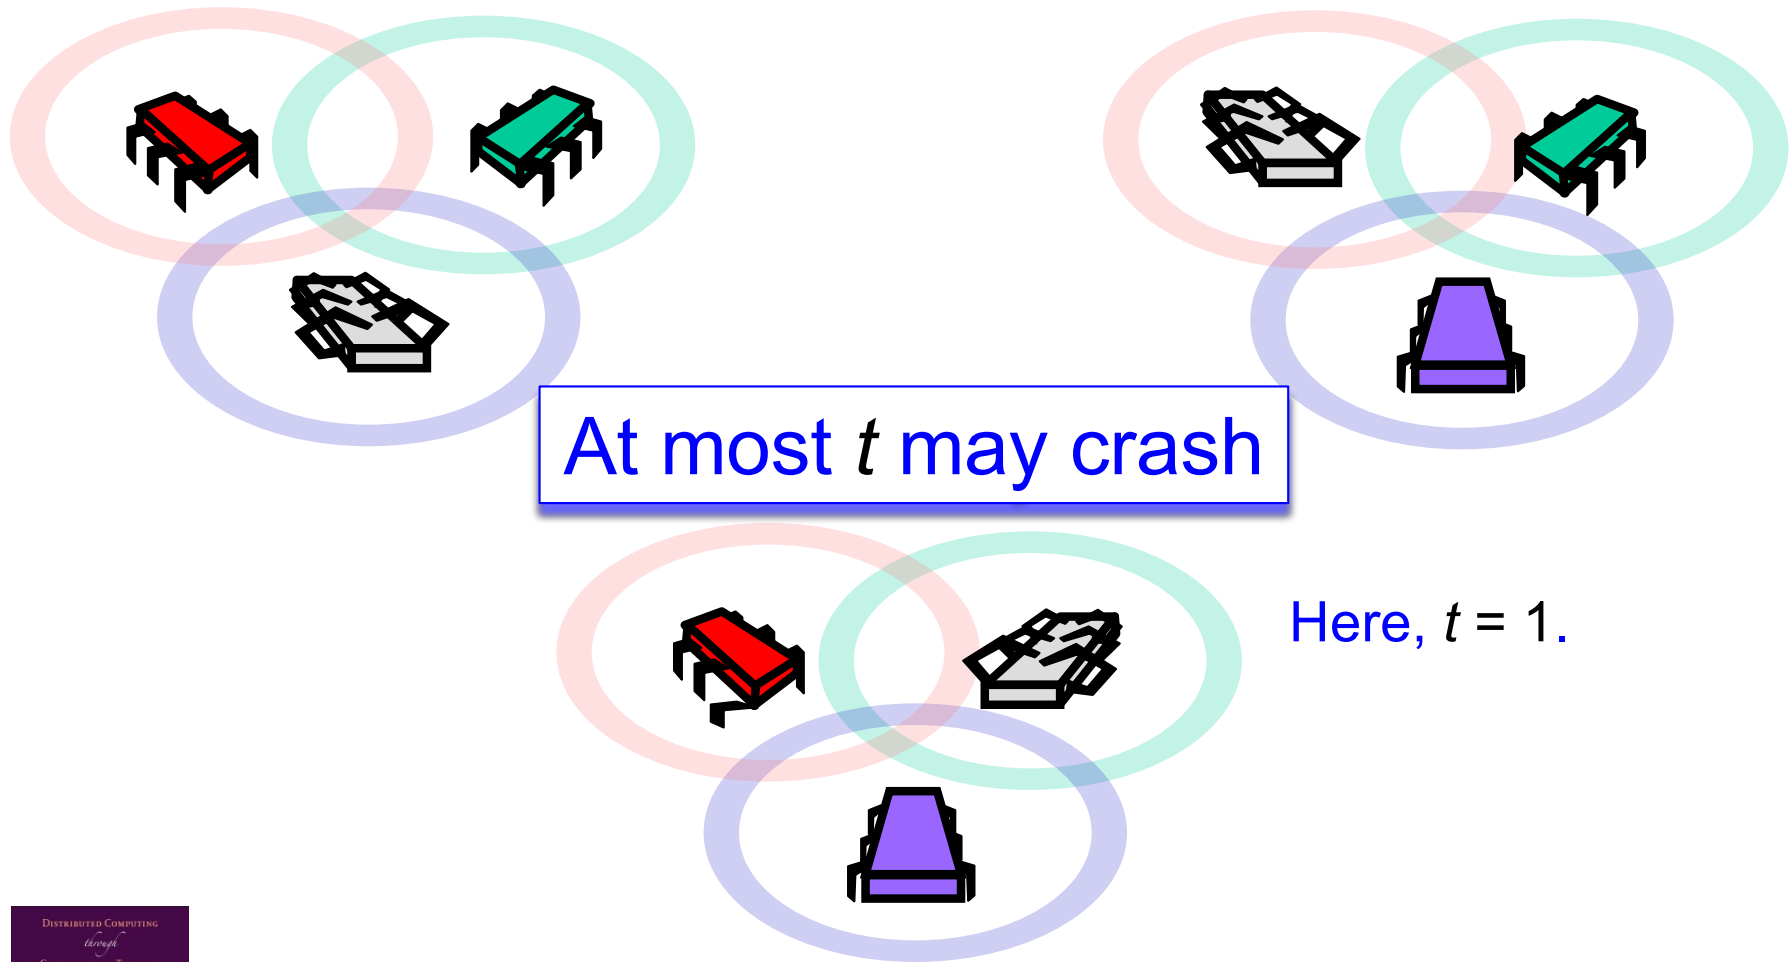
Which ones?



Wait-Free Failure Model



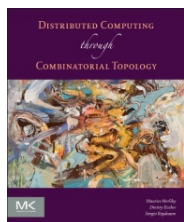
t -Resilient Failure Model



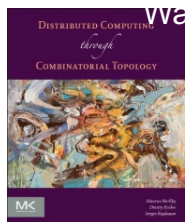
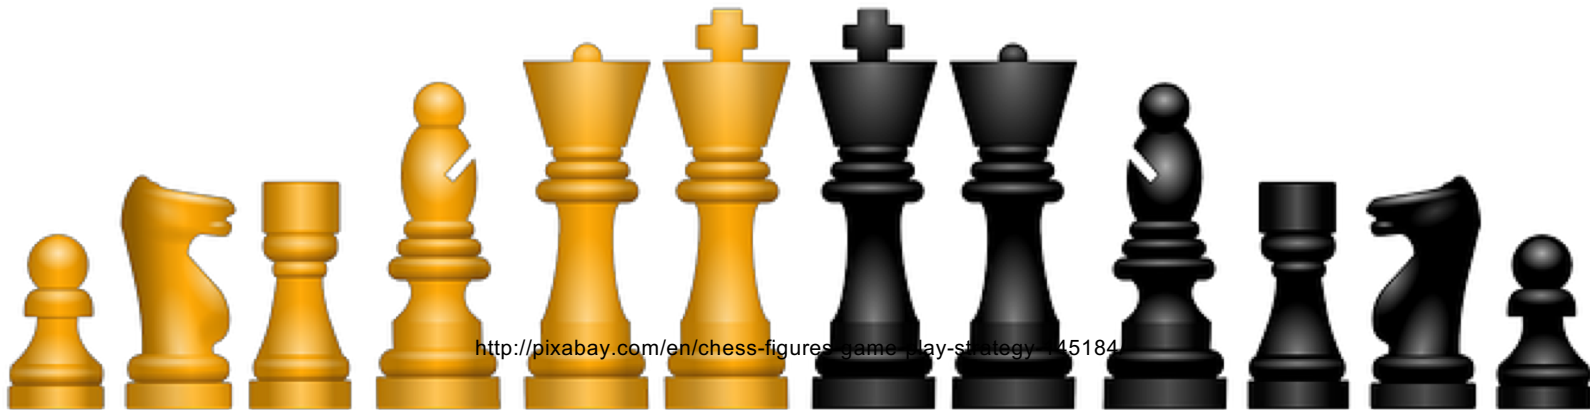
Correlated Failures



Processes on same server may crash

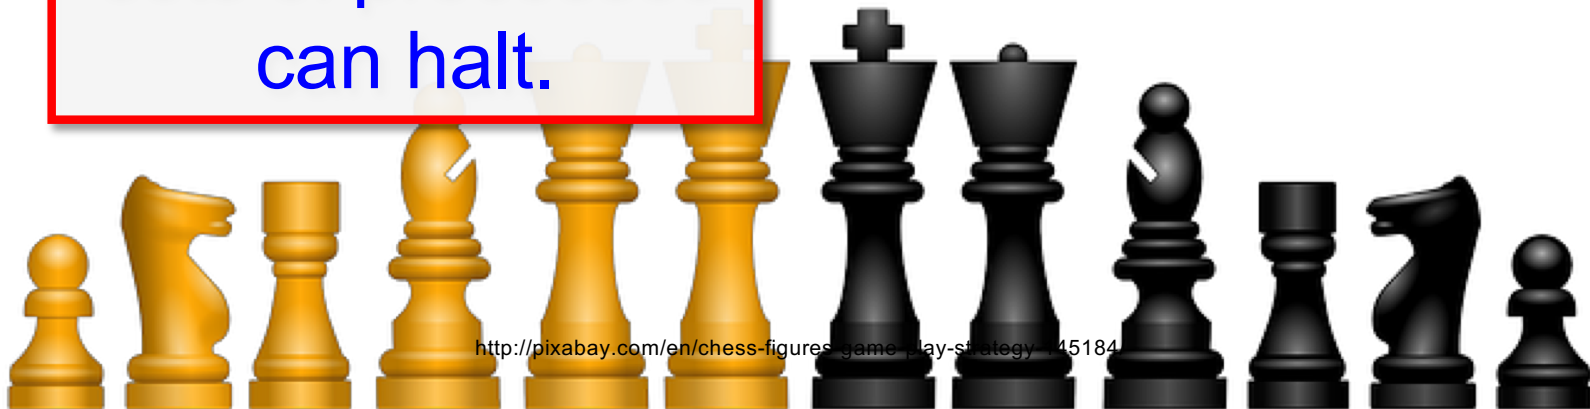


Adversaries

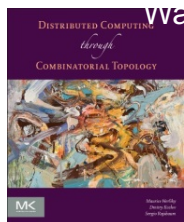


Adversaries

Determine which
sets of processes
can halt.

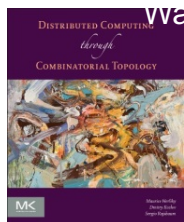
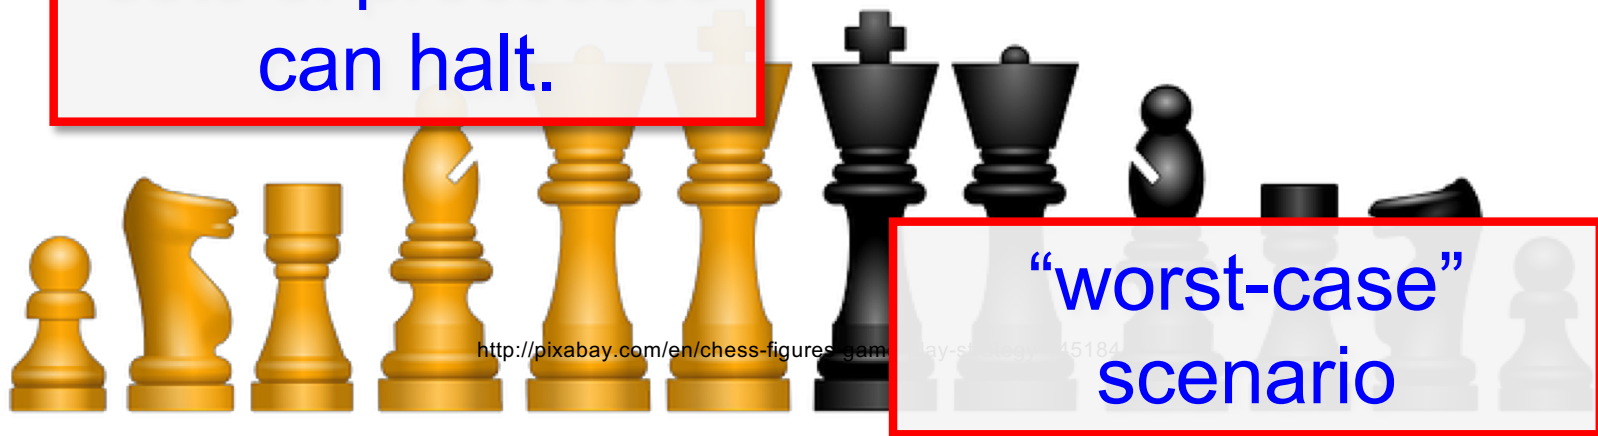


<http://pixabay.com/en/chess-figures-game-play-strategy-45184/>

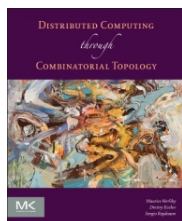


Adversaries

Determine which
sets of processes
can halt.



Timing Models

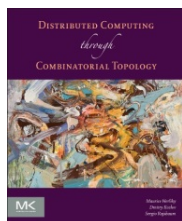


Distributed Computing through
Combinatorial Topology

Timing Models



Processes share a clock

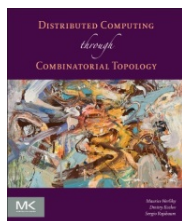


Timing Models



Processes share a clock

Synchronous



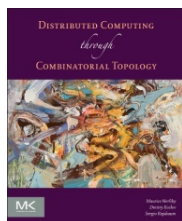
Timing Models



Processes share a clock

Synchronous

Processes do not share a clock



Timing Models

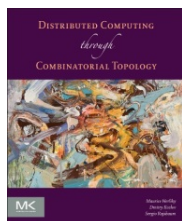


Processes share a clock

Synchronous

Processes do not share a clock

Asynchronous



Timing Models



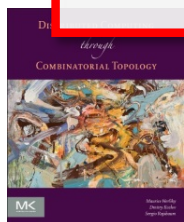
Processes share a clock

Synchronous

Processes do not share a clock

Asynchronous

Processes have approximately-synchronized clocks



Timing Models



Processes share a clock

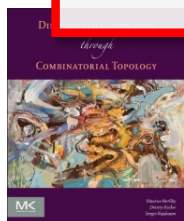
Synchronous

Processes do not share a clock

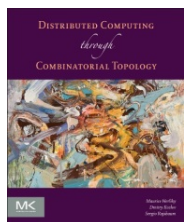
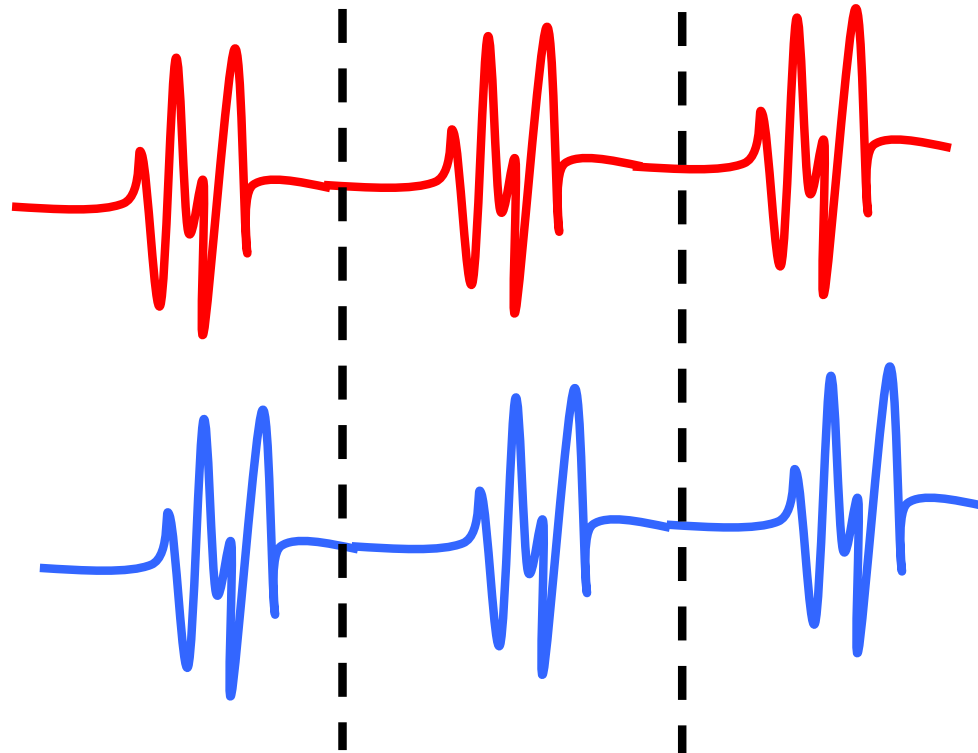
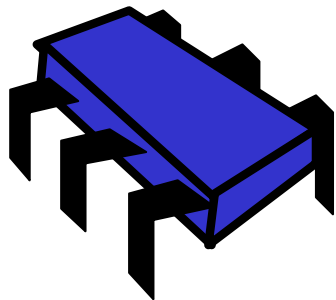
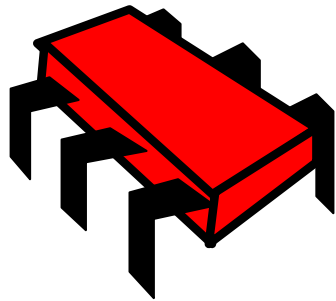
Asynchronous

Processes have approximately-synchronized clocks

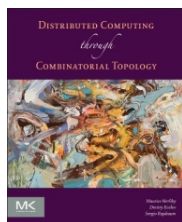
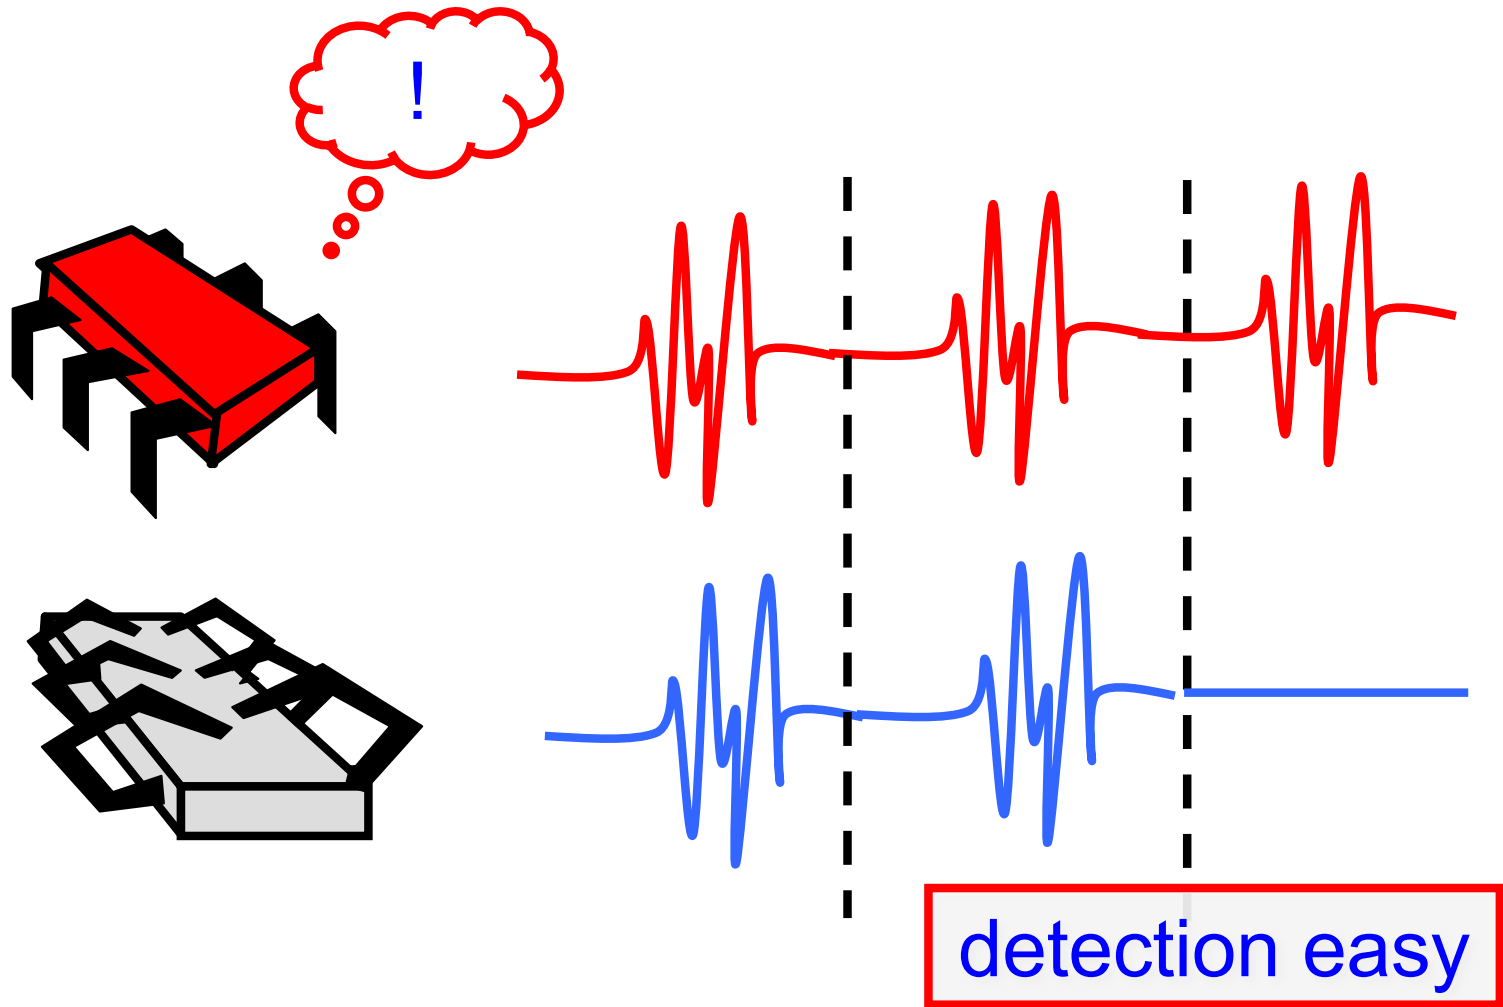
Semi-synchronous



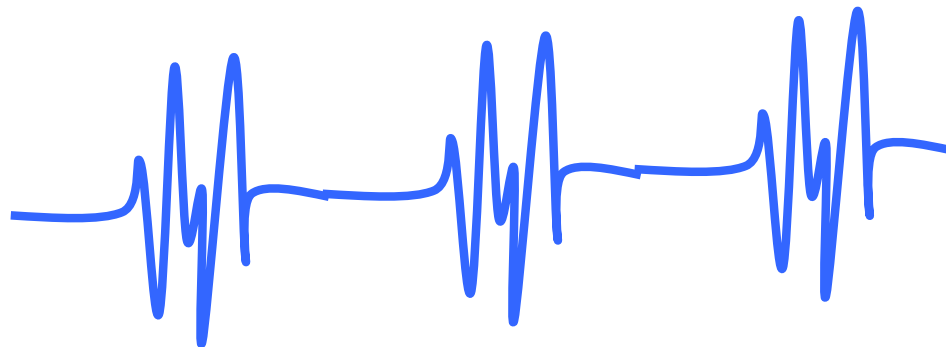
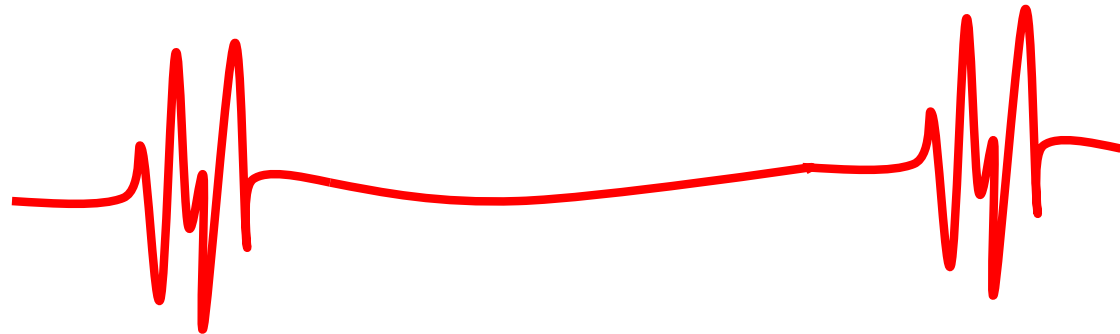
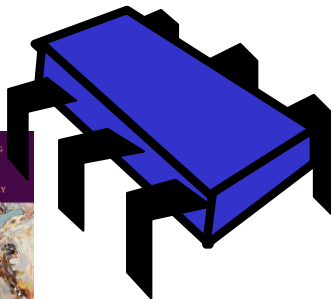
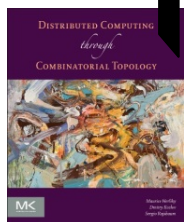
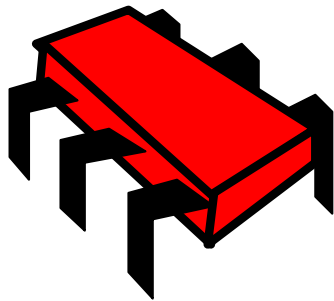
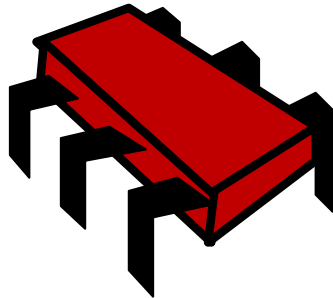
Synchronous



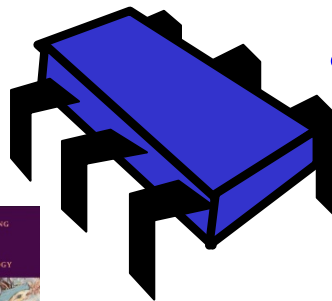
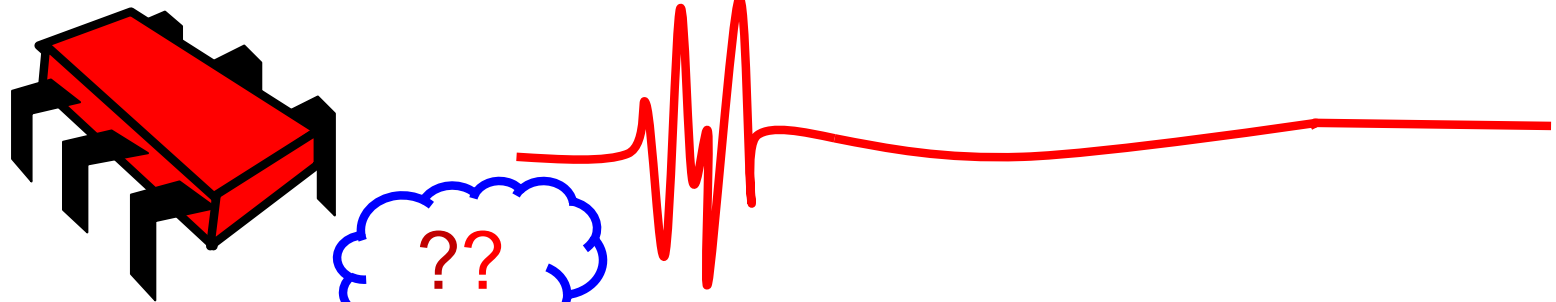
Synchronous Failures



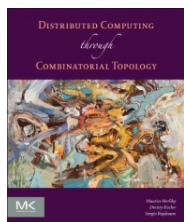
Asynchronous



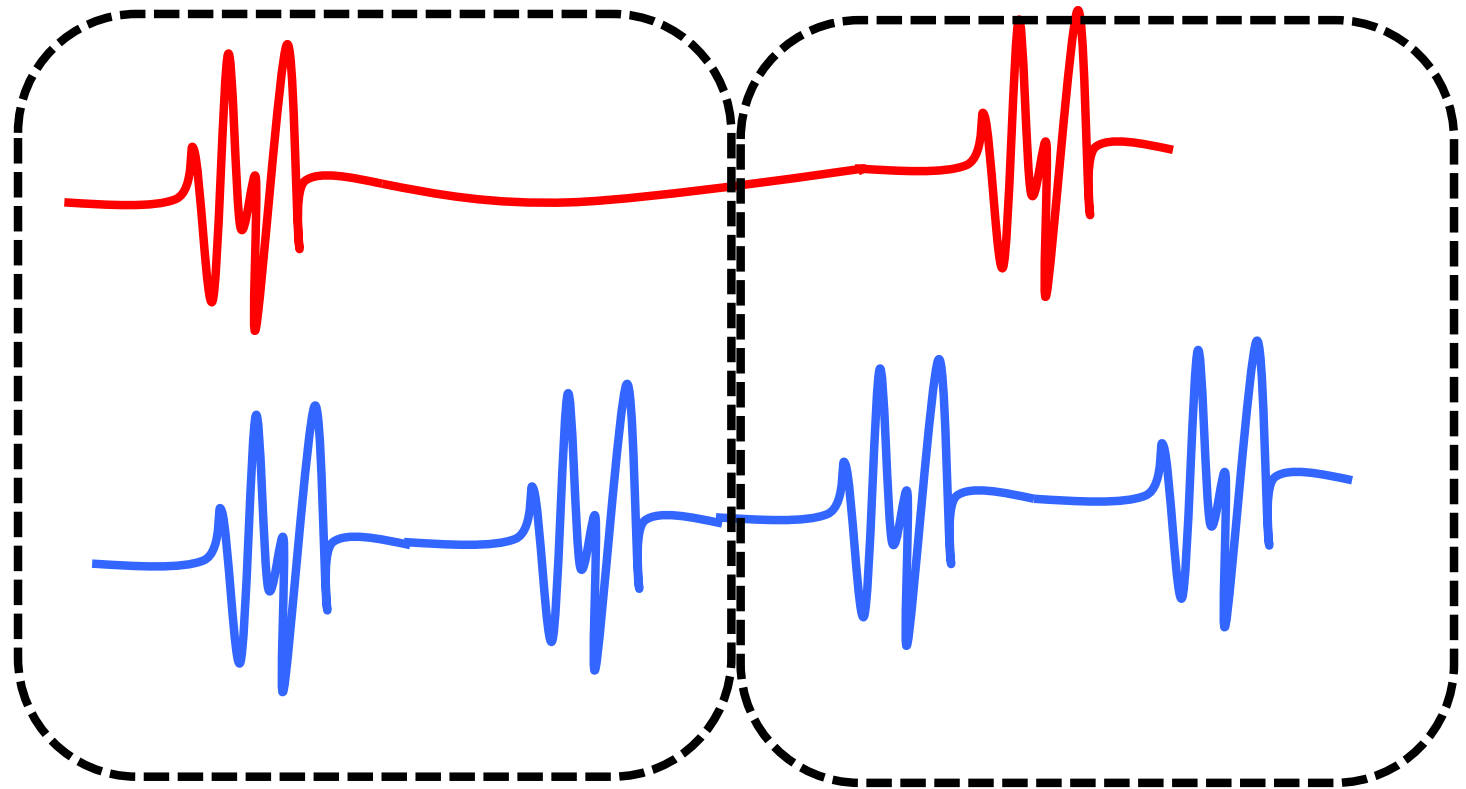
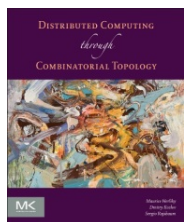
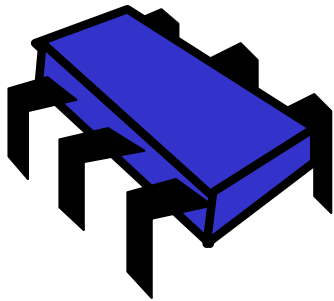
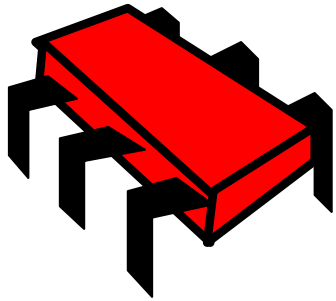
Asynchronous Failures



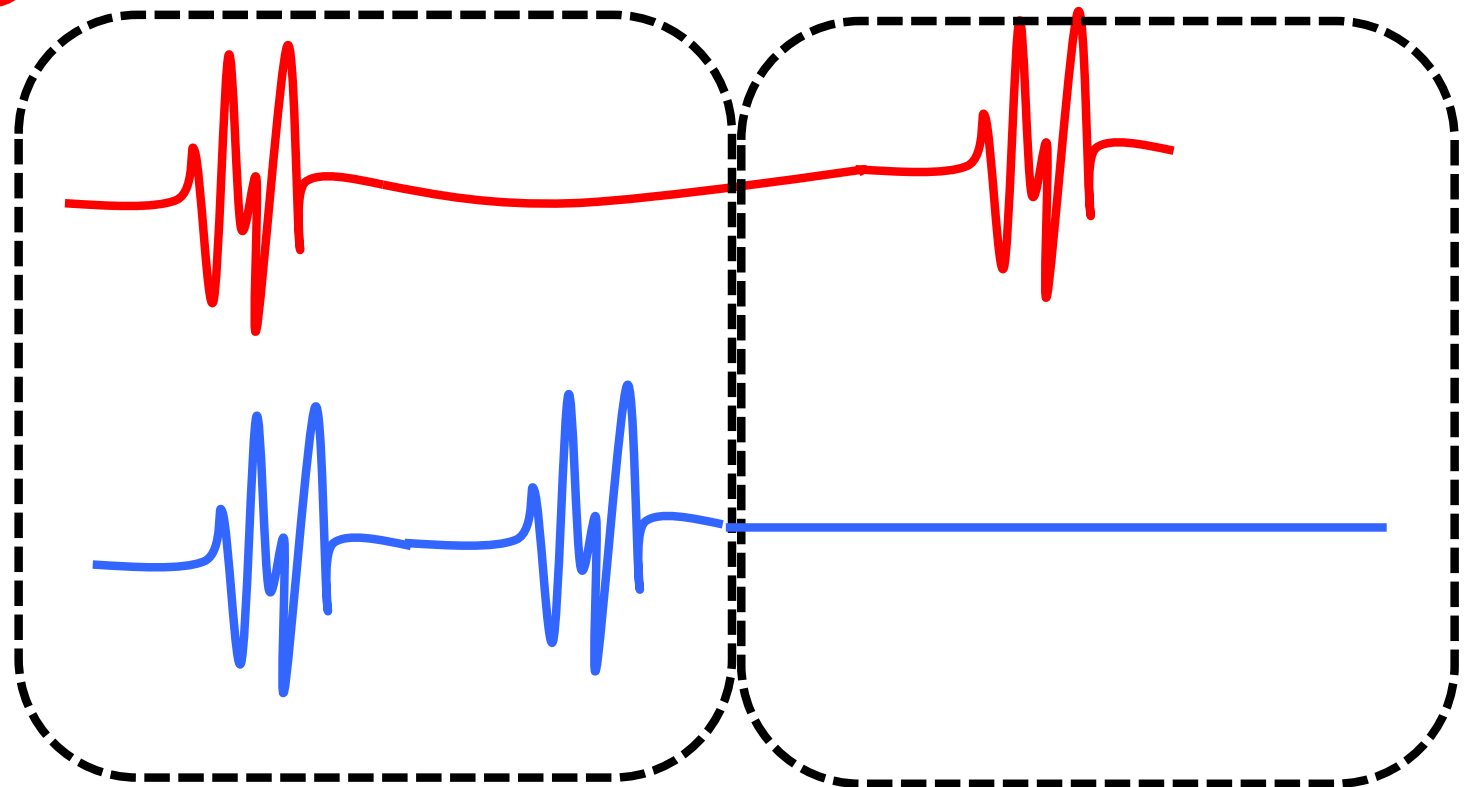
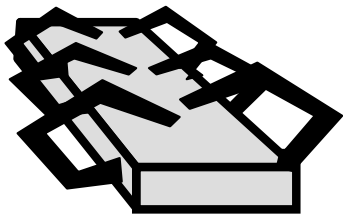
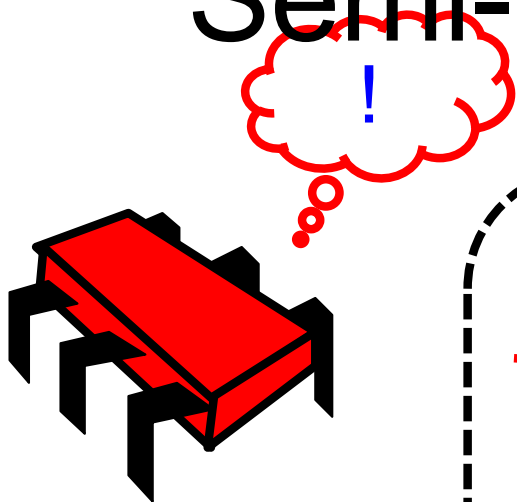
detection impossible



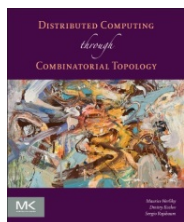
Semi-Synchronous



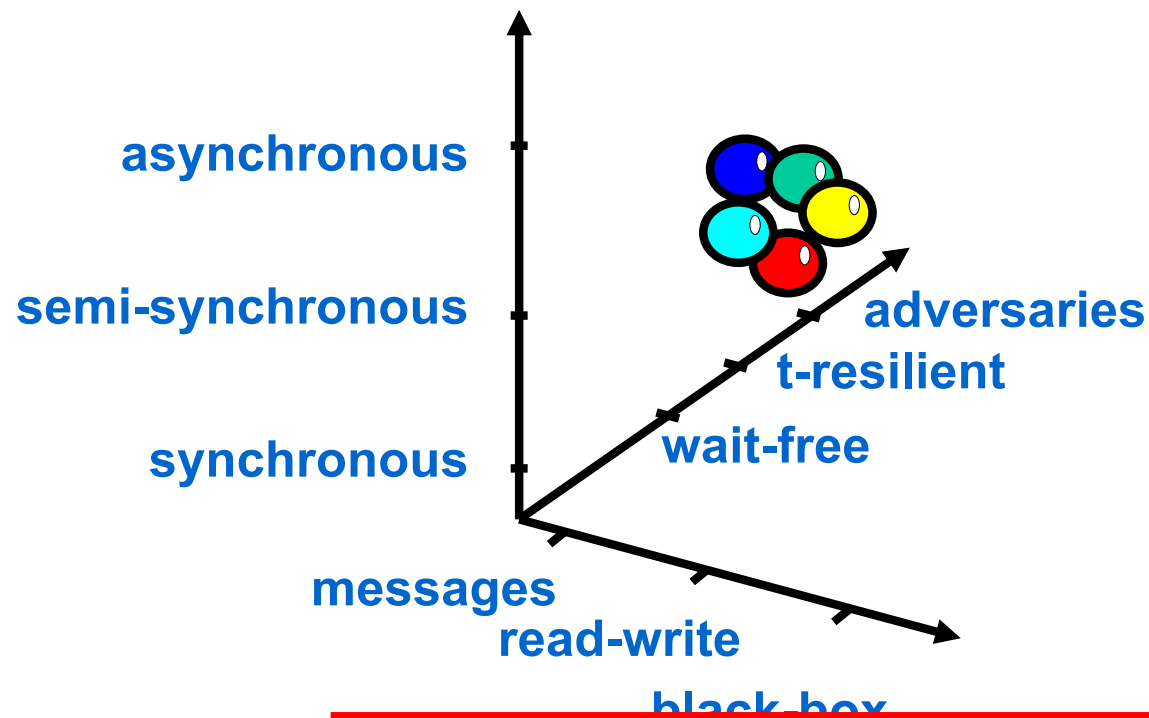
Semi-Synchronous Failures



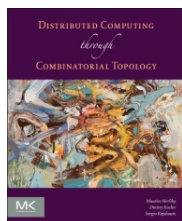
detection slow



Computation Model Space

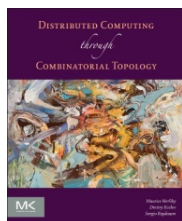
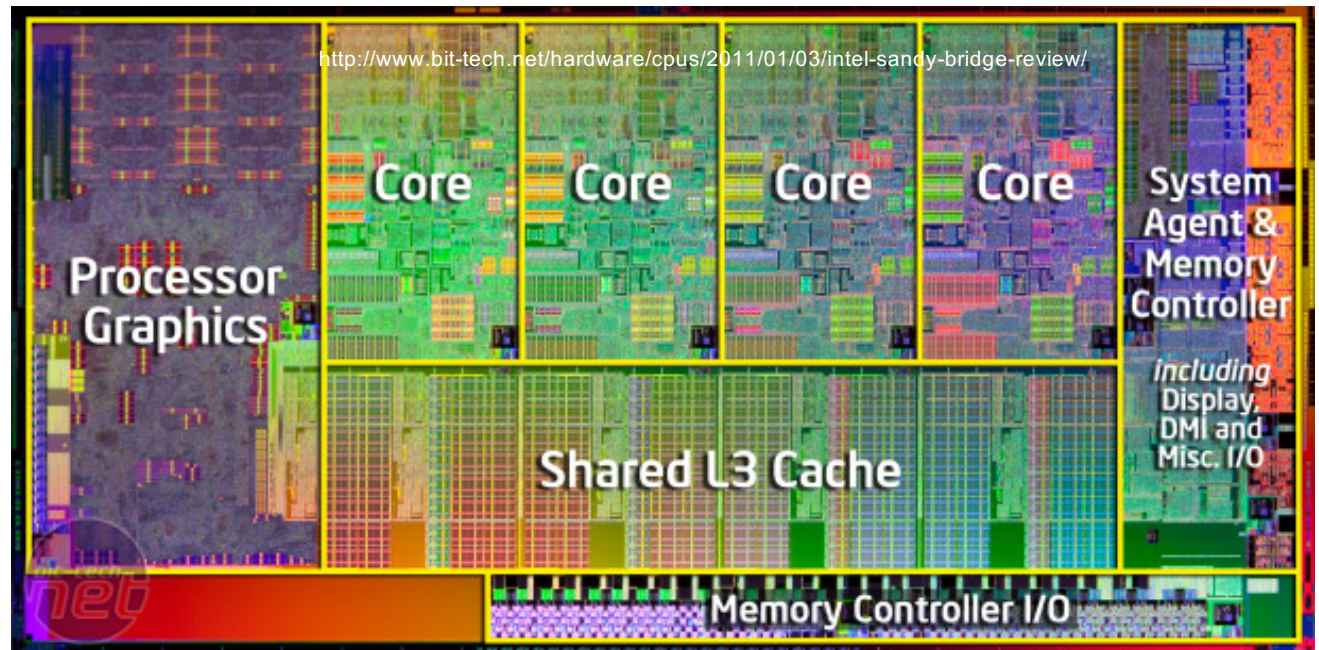


Which combinations make sense?



Multicores

Asynchronous
Wait-free
Shared Memory

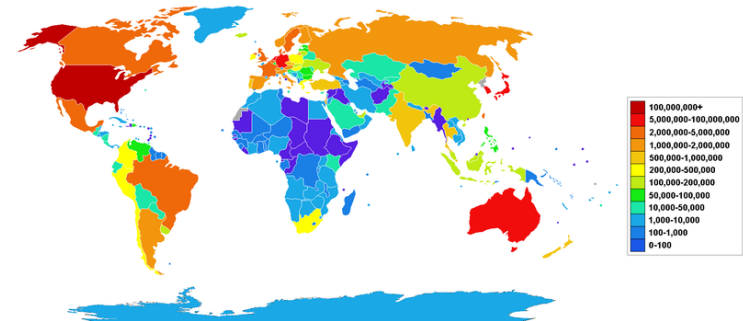


Distributed Computing

Asynchronous
Message-passing

Internet

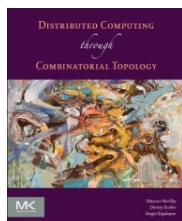
Sensor
network



http://commons.wikimedia.org/wiki/File:Internet_hosts.PNG



http://commons.wikimedia.org/wiki/File:XBee_Series_2_with_Whip_Antenna.jpg



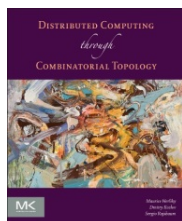
Parallel Computing

Synchronous
Message-passing
(or shared memory)

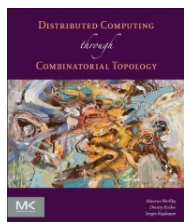
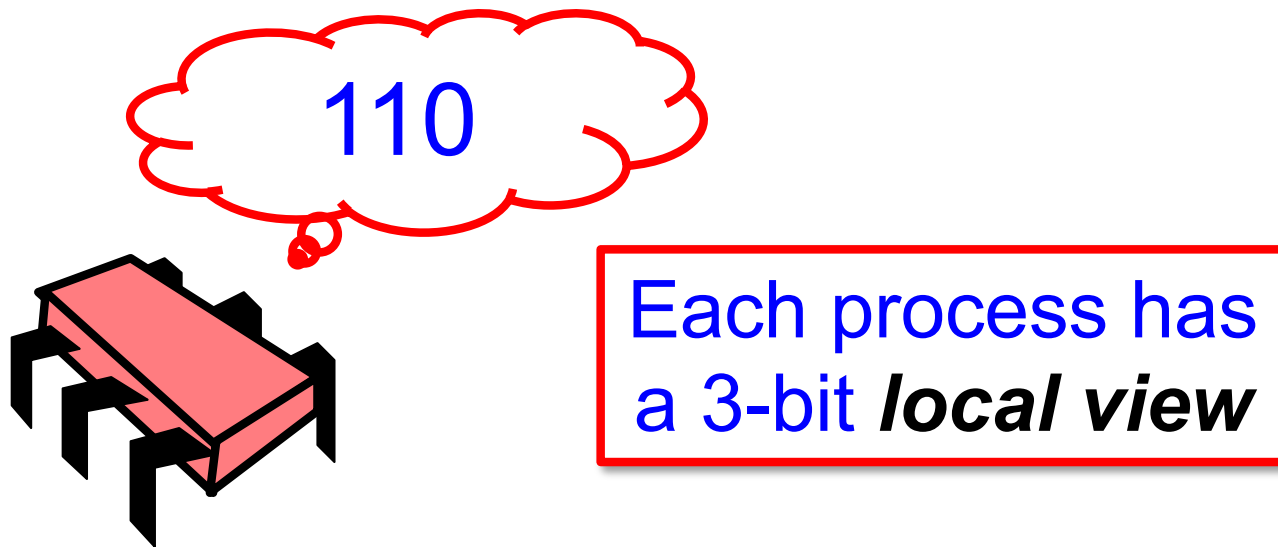


http://commons.wikimedia.org/wiki/File:GeForce_2_mx400_gpu.JPG

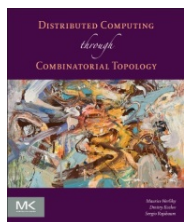
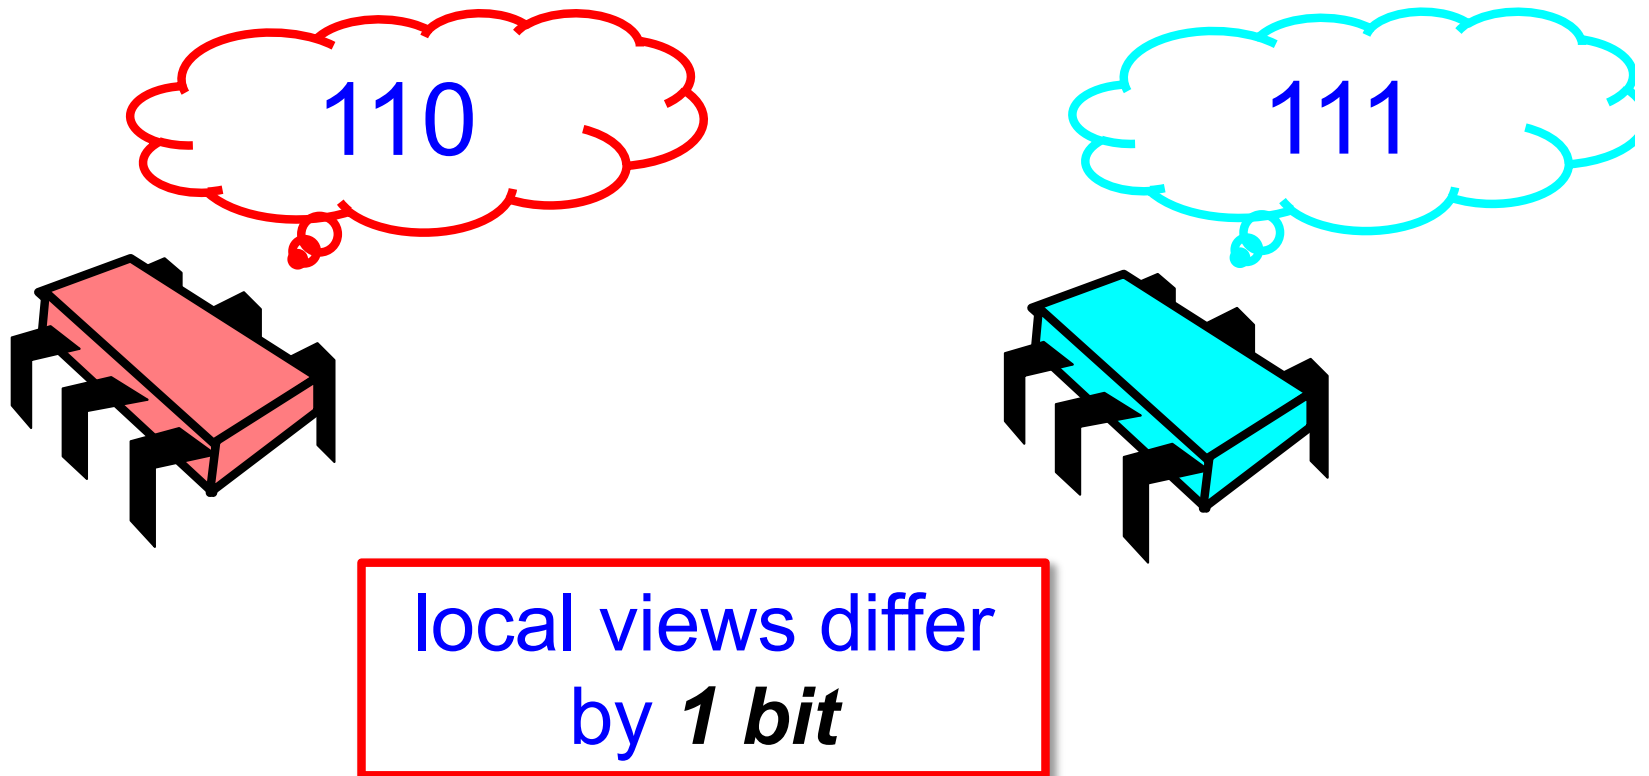
GPU



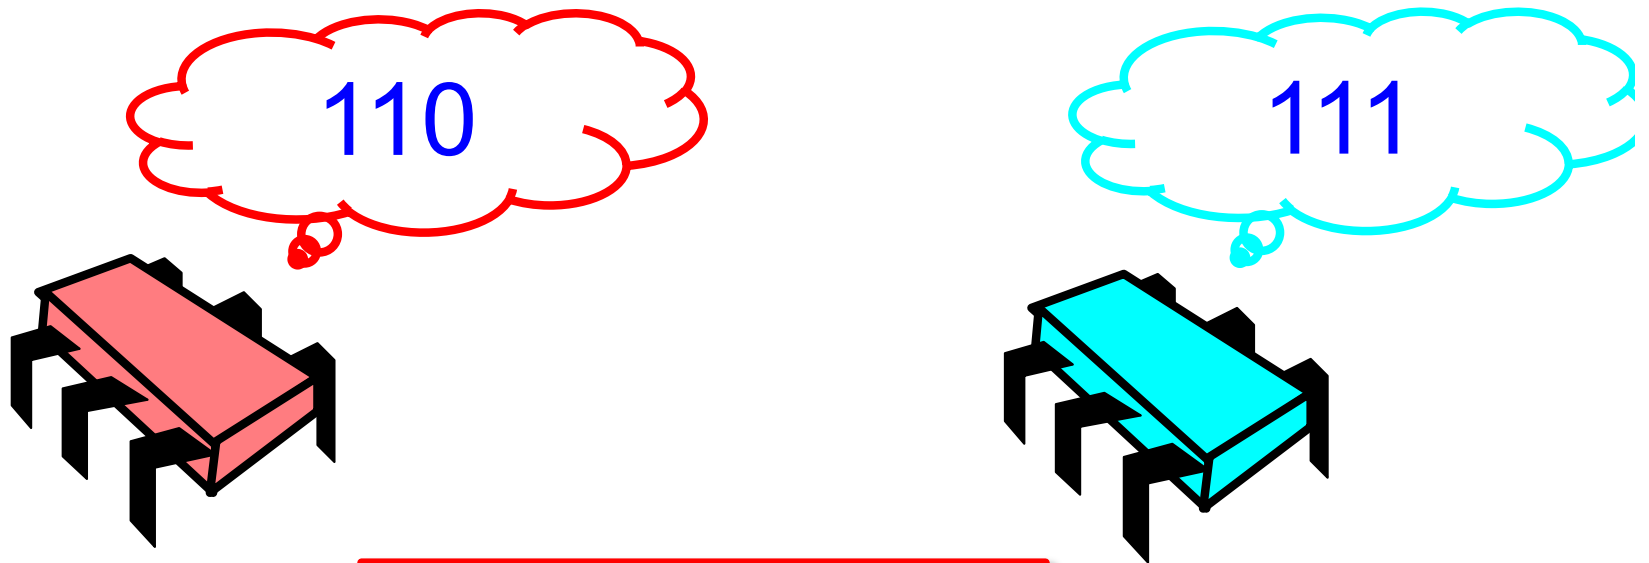
Local Views



Multiple Local Views

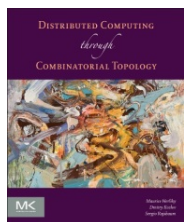


Multiple Local Views

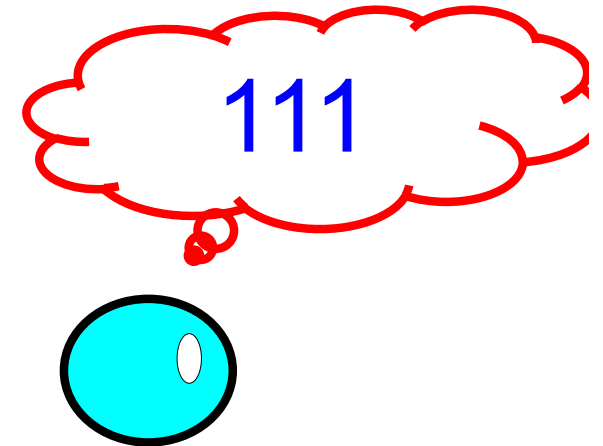
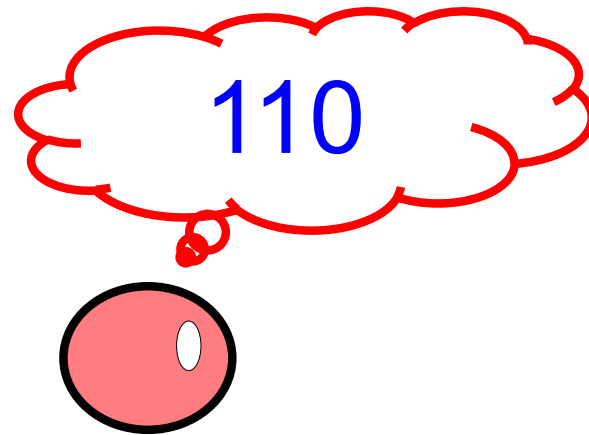


local views differ
by **1 bit**

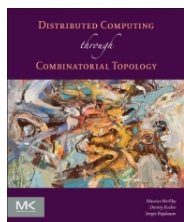
but no process
knows which one



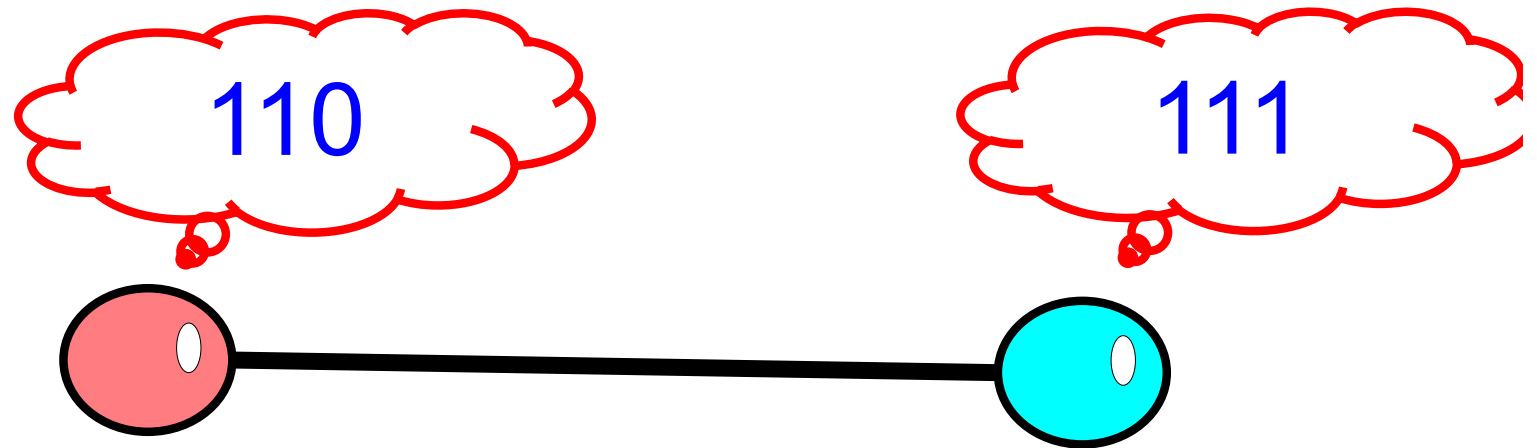
Multiple Local Views



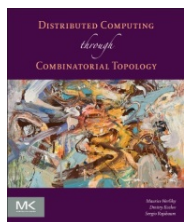
each view is
represented by a
labeled vertex

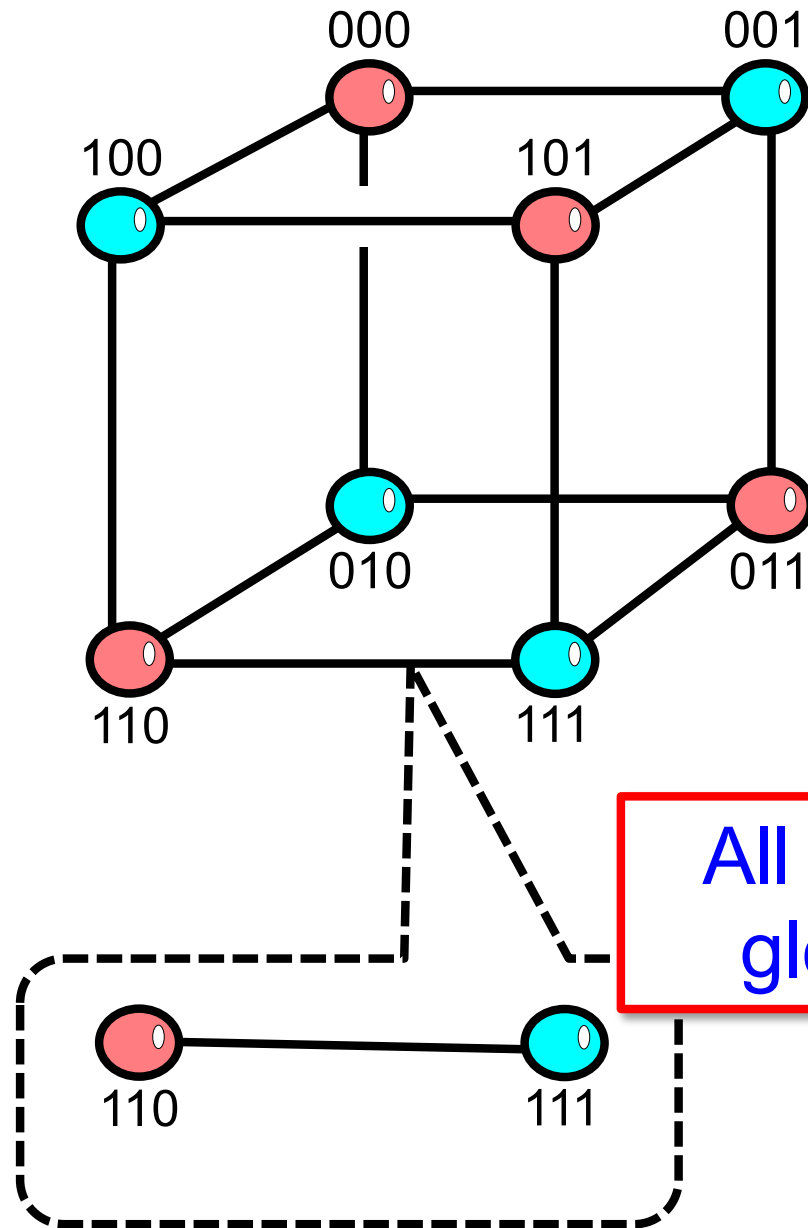


Global States

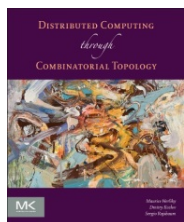


compatible views
represented by an
edge



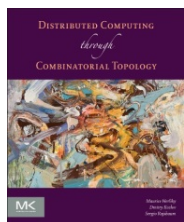
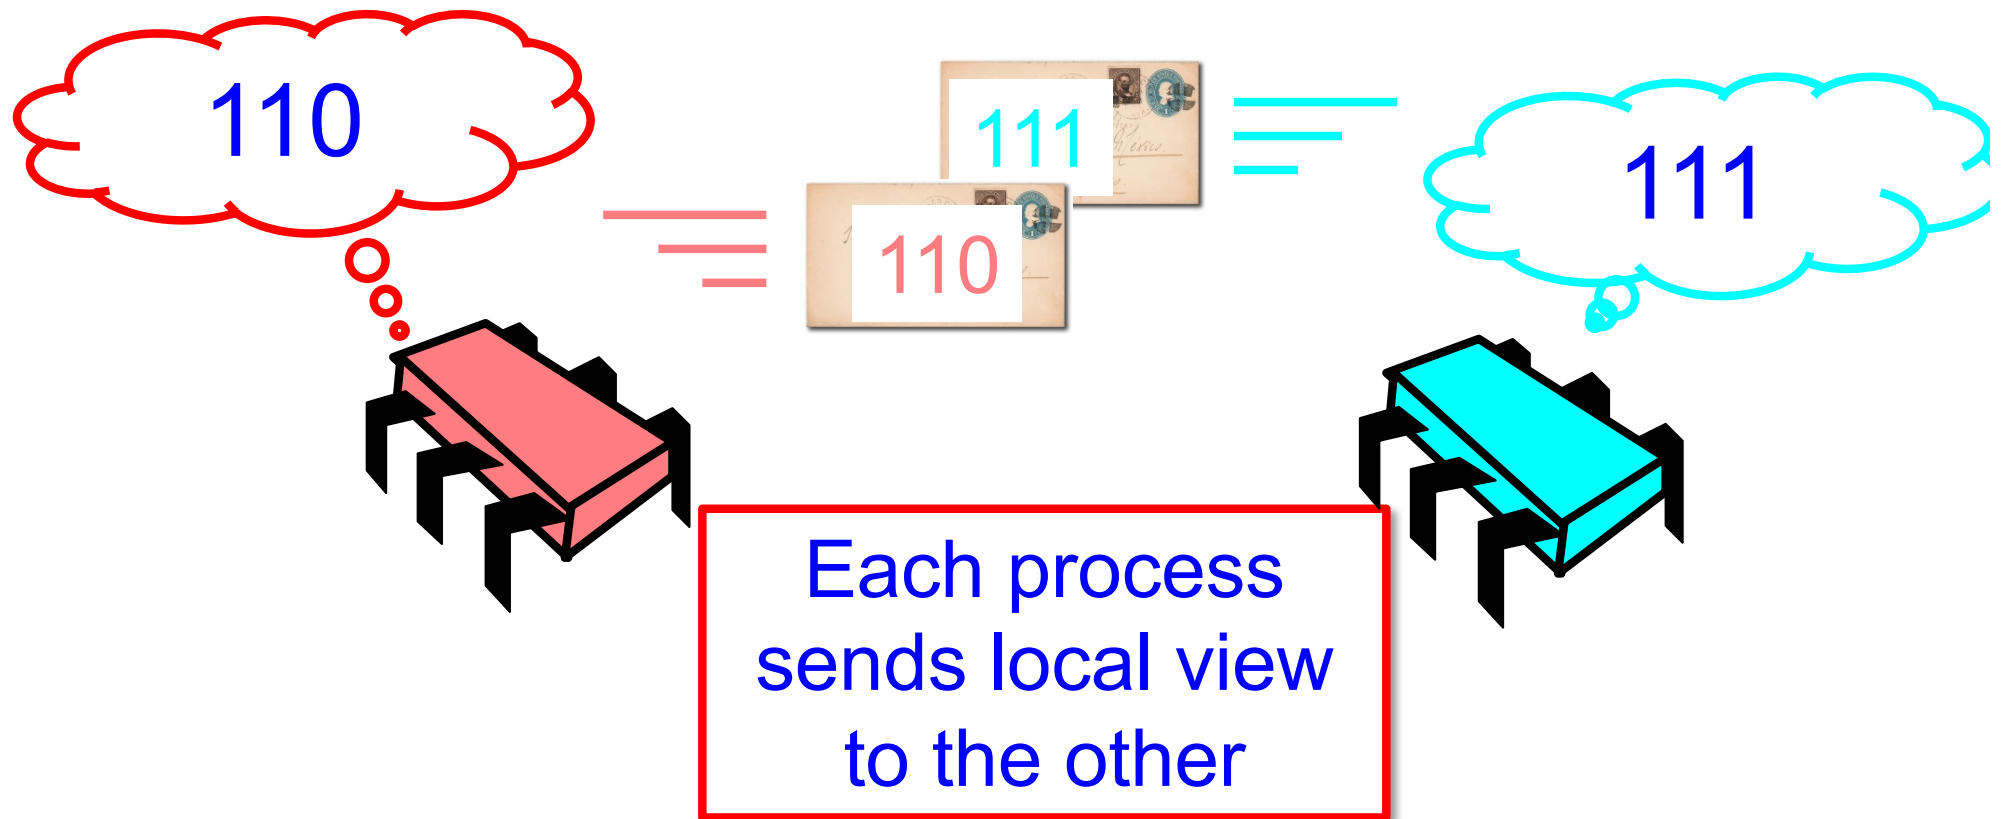


All (?) possible
global states

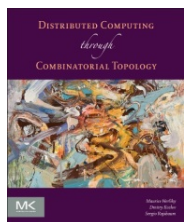
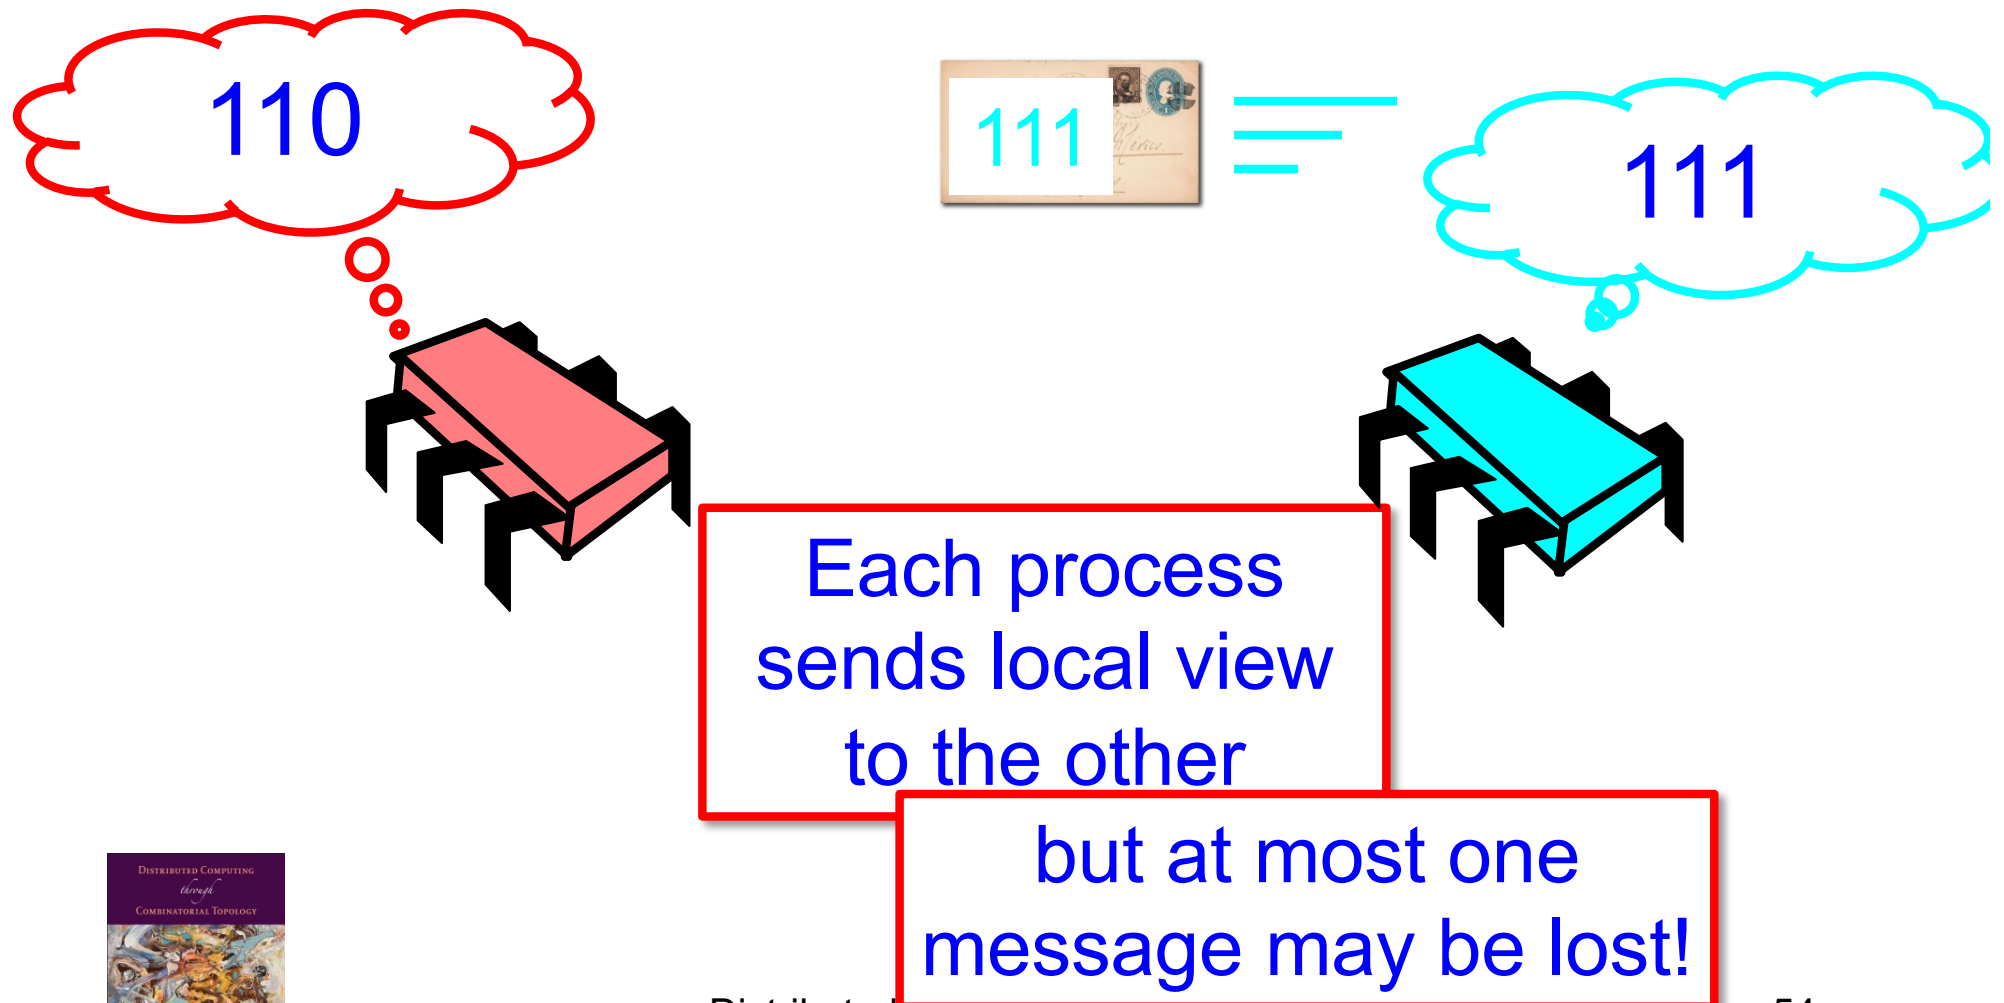


Distributed Computing through
Combinatorial Topology

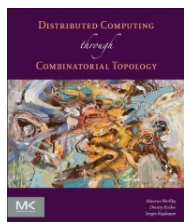
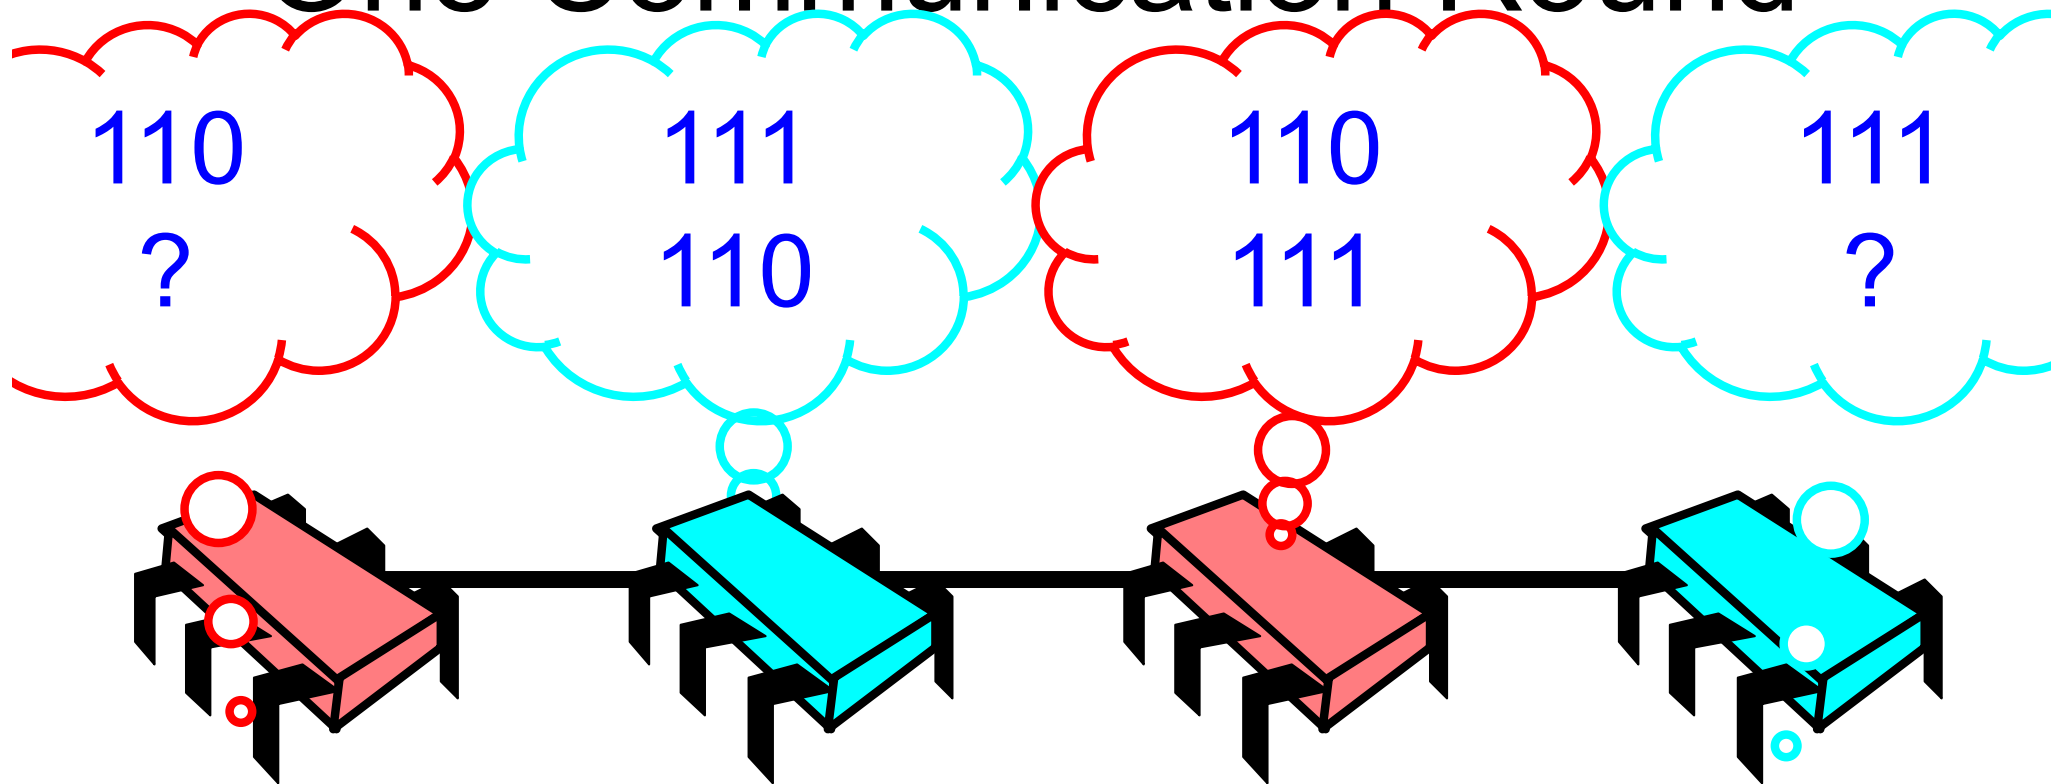
Communication



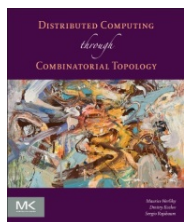
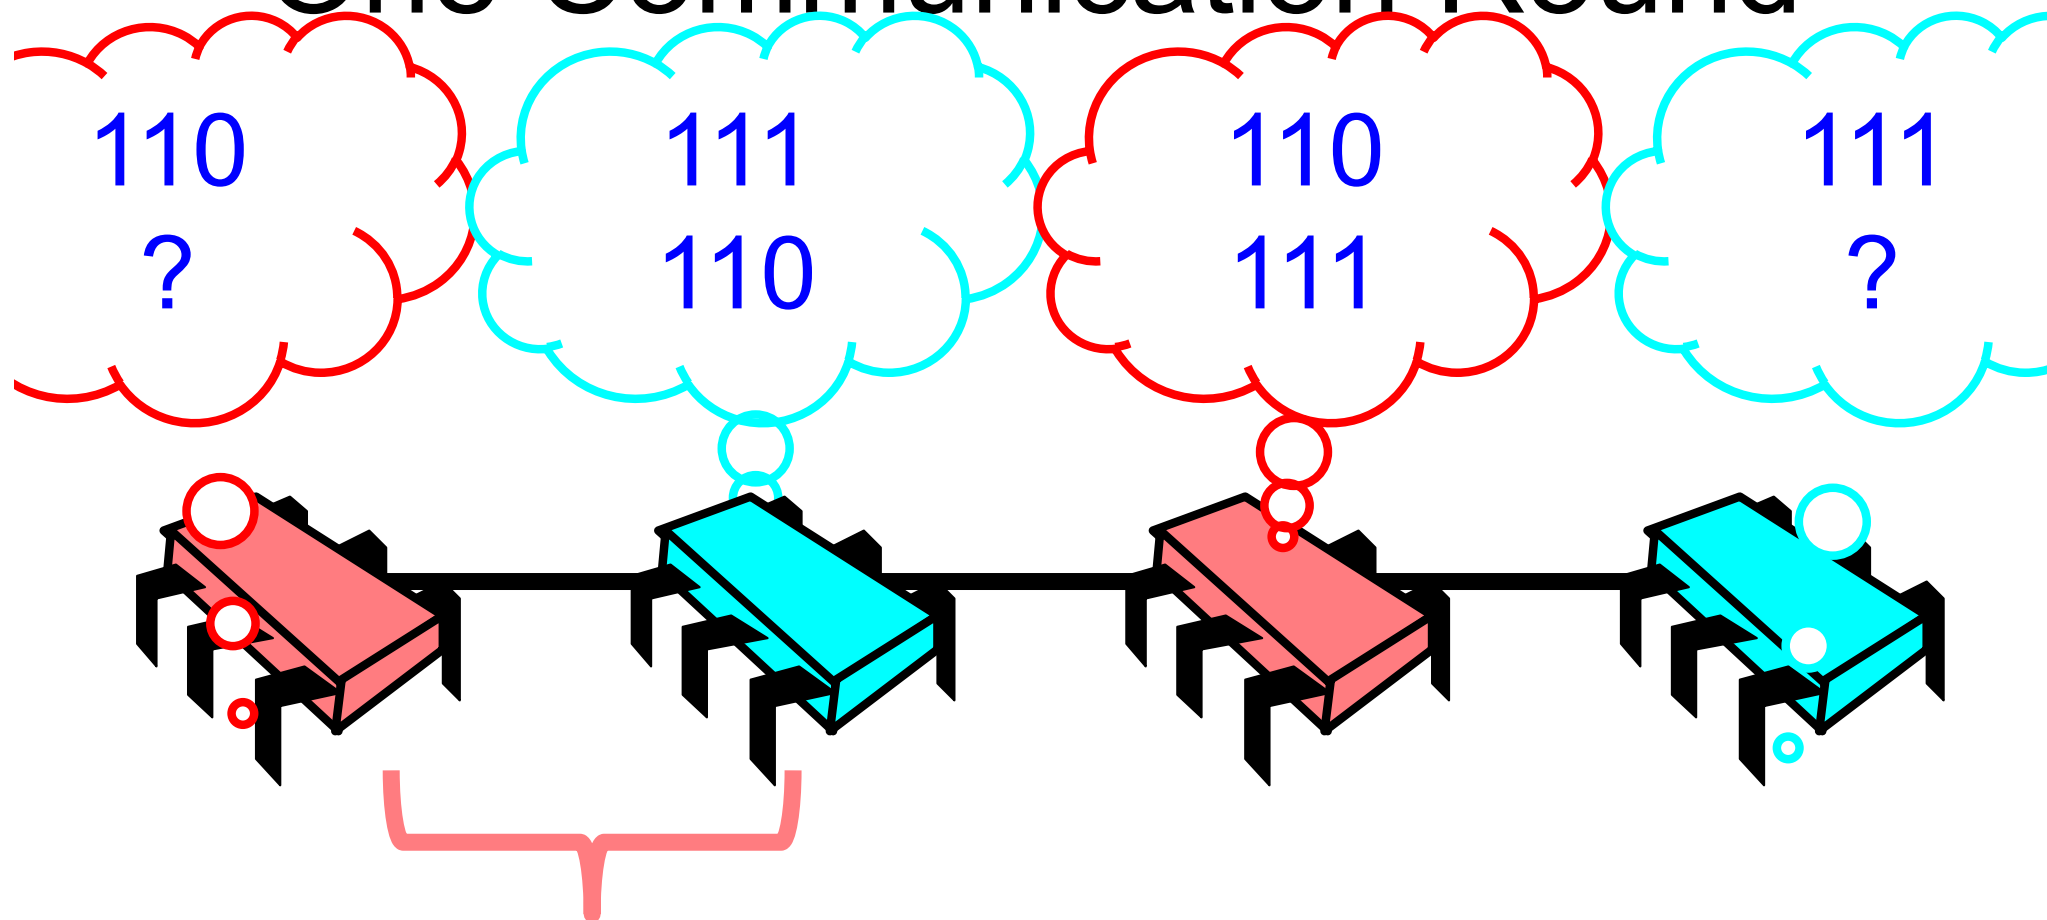
Communication



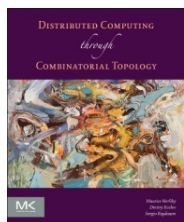
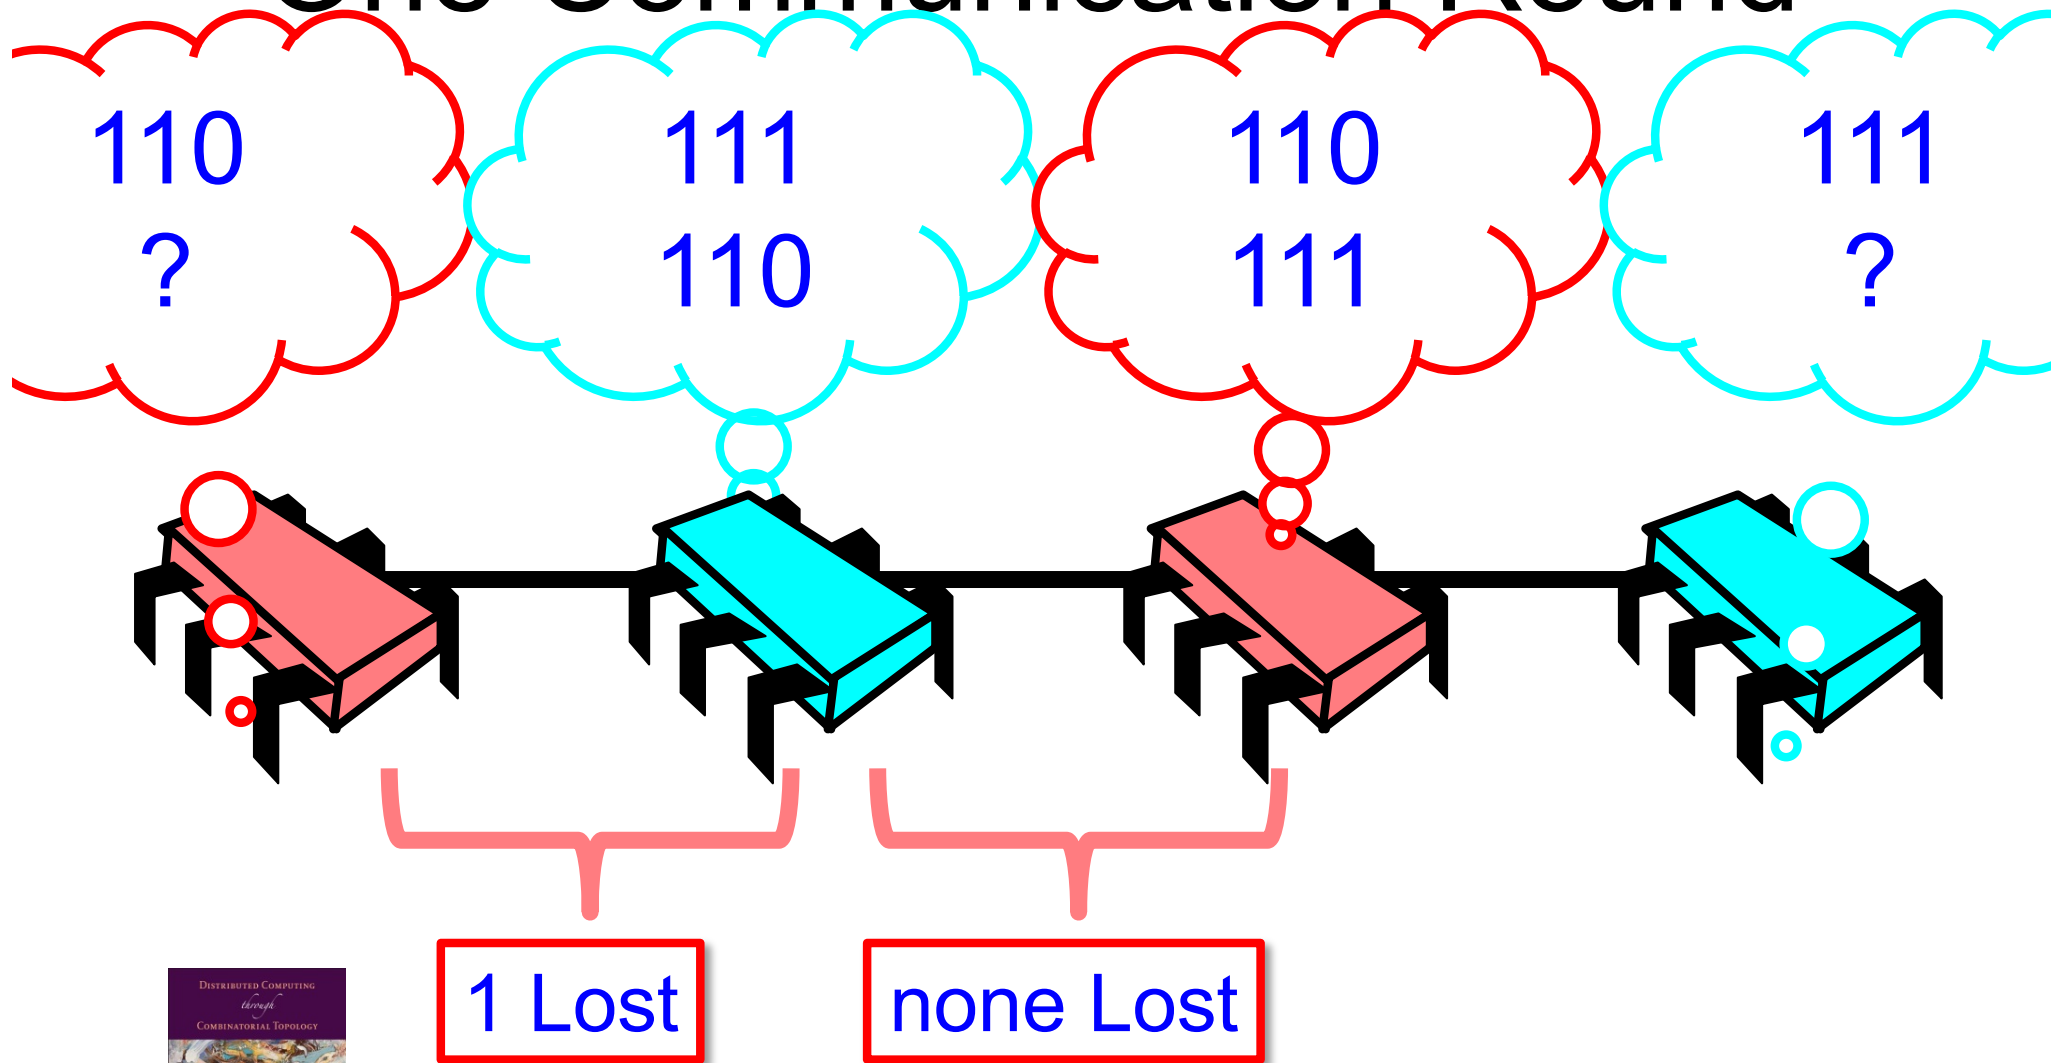
One Communication Round



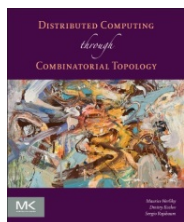
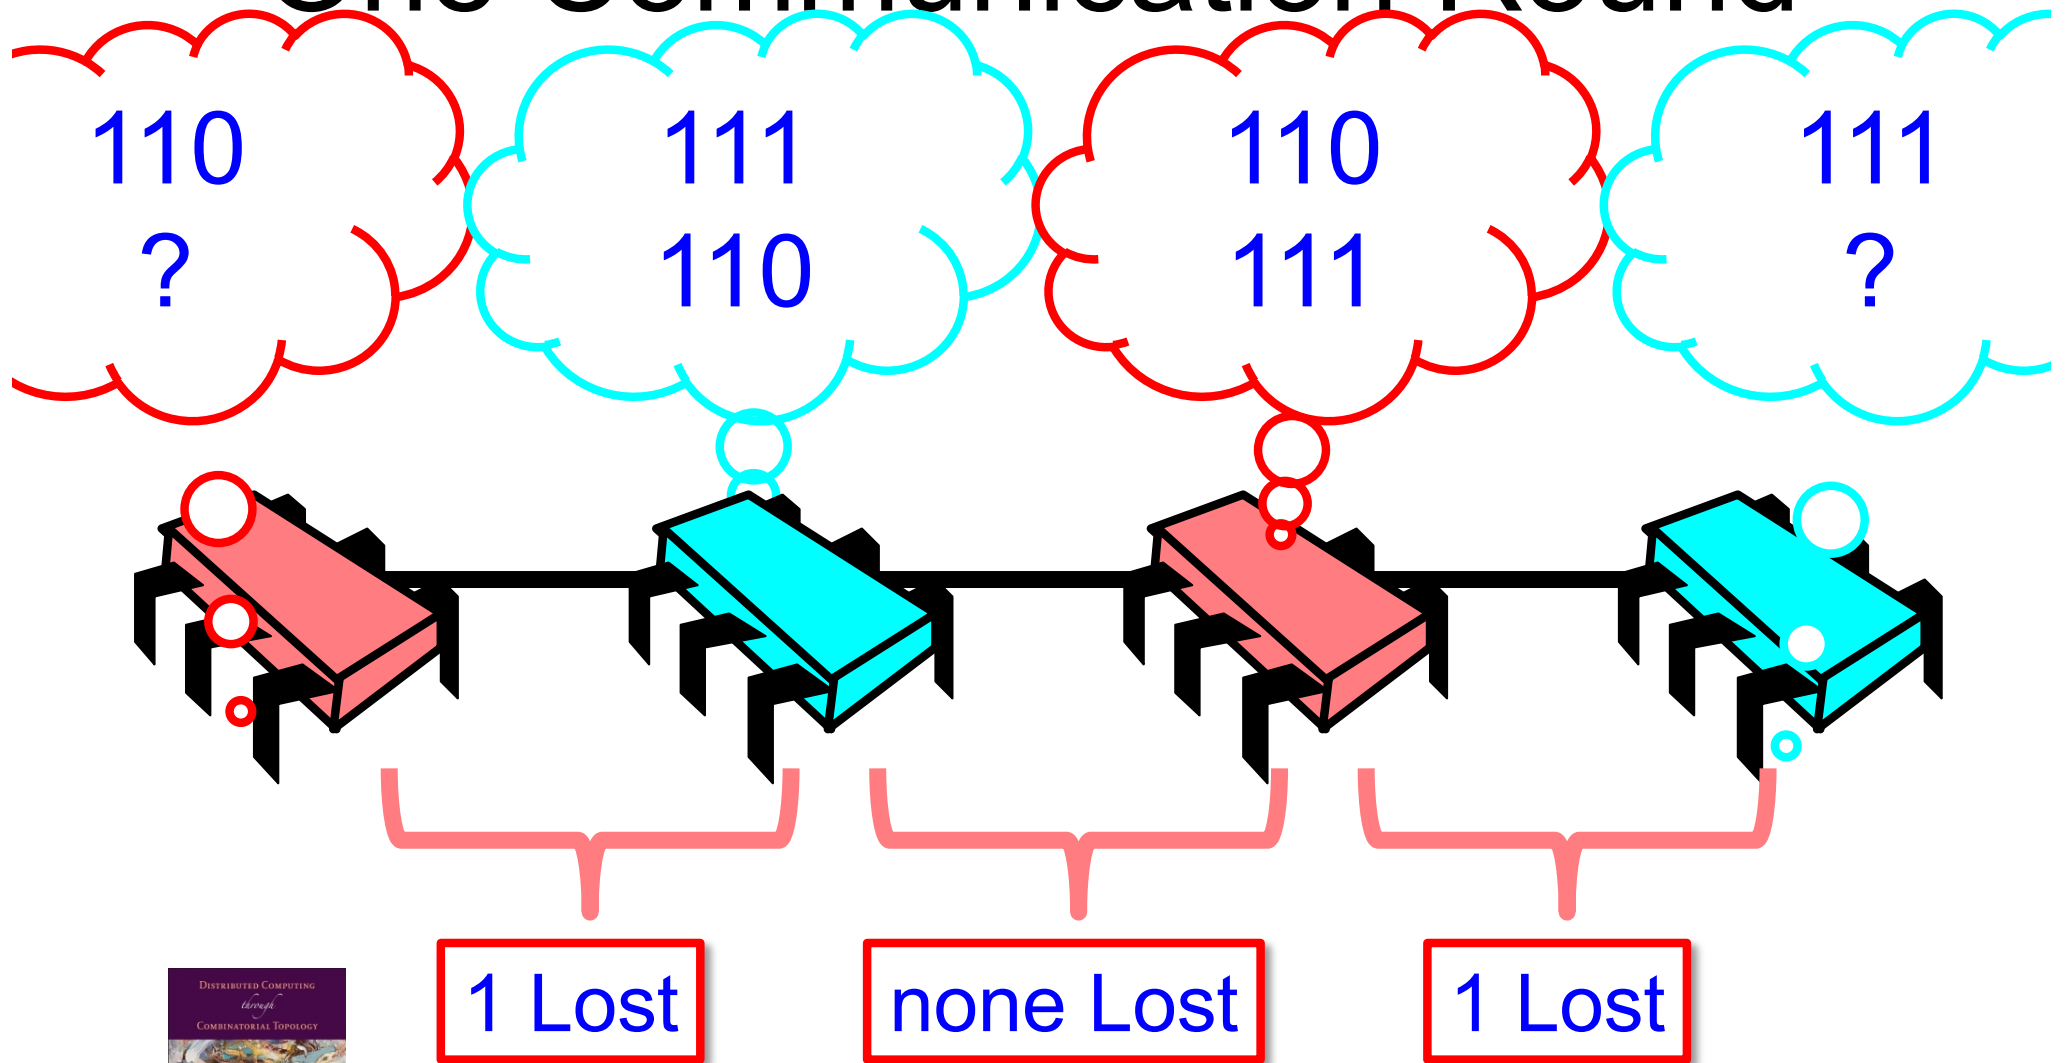
One Communication Round



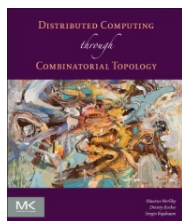
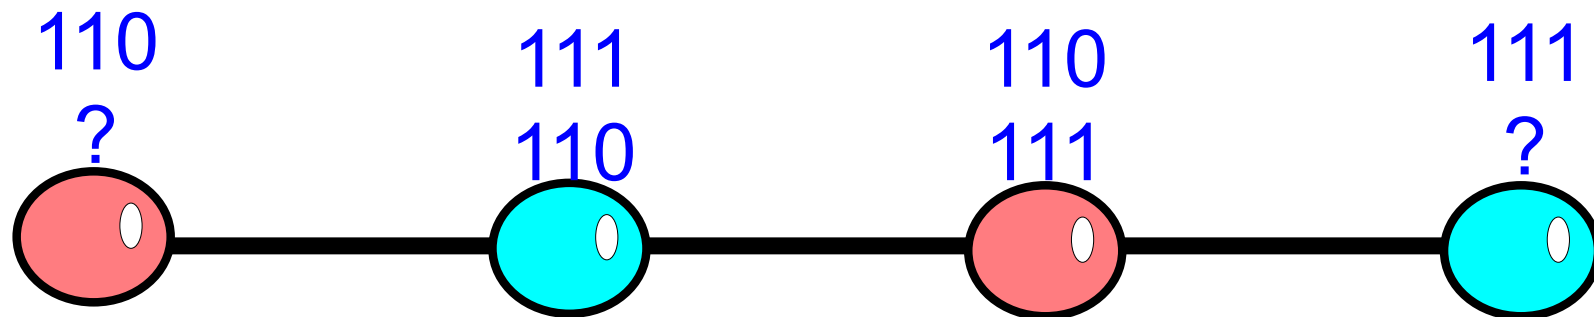
One Communication Round



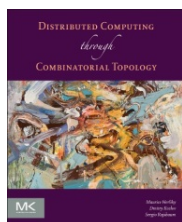
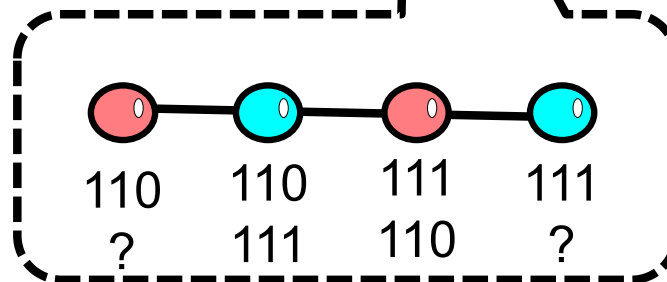
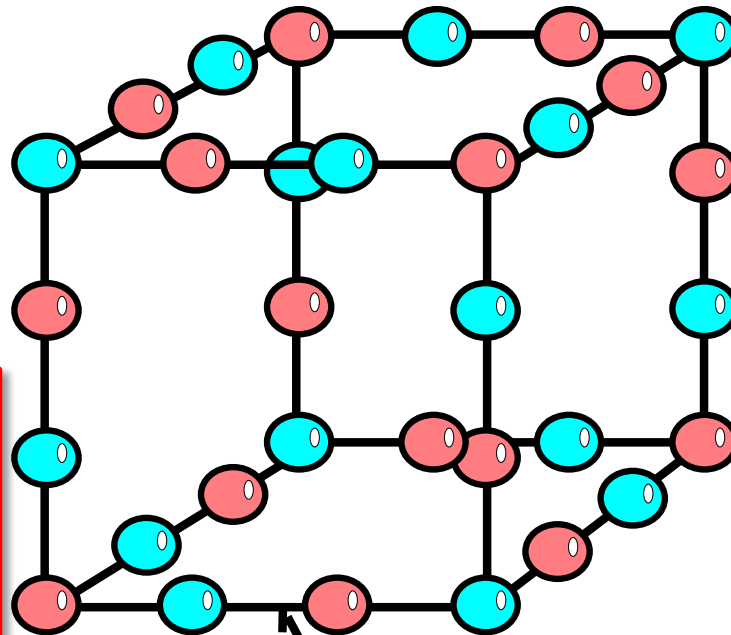
One Communication Round



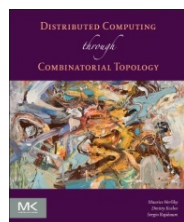
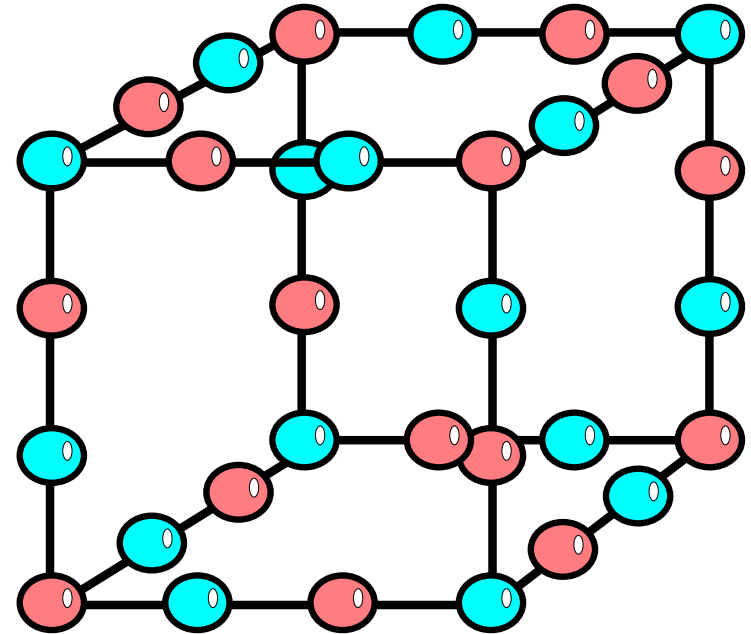
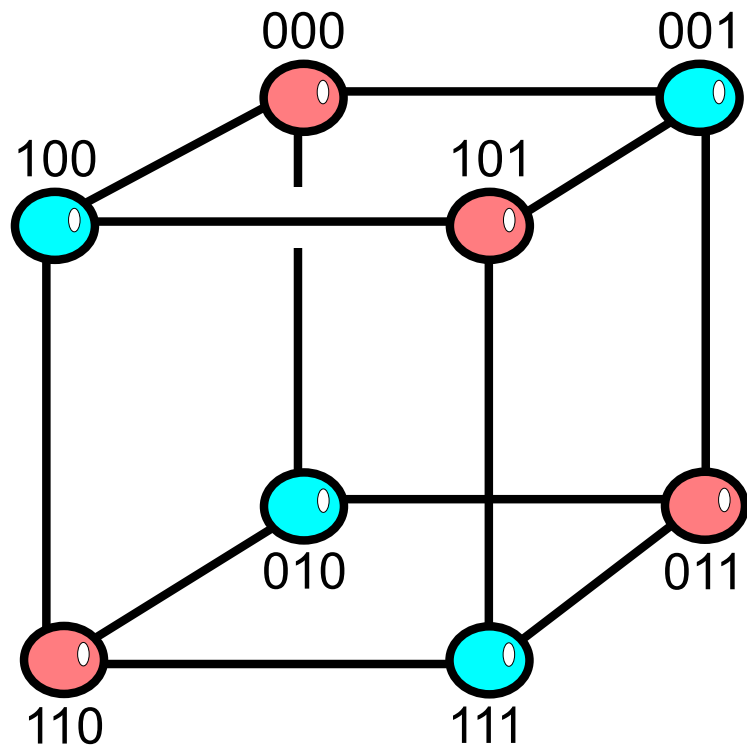
One Communication Round

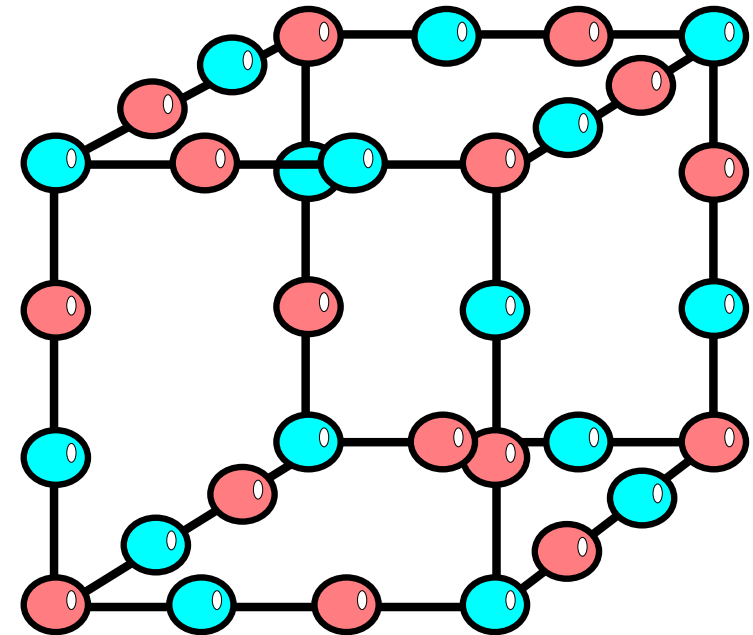
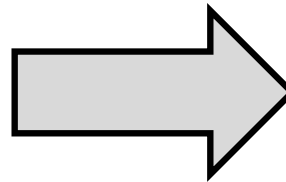
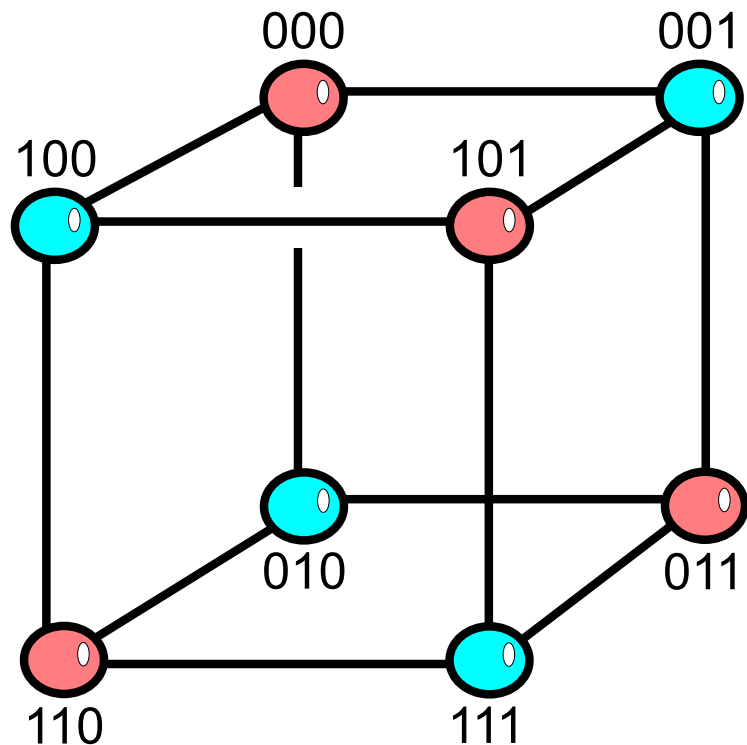


All possible global
states after one
round unreliable
communication



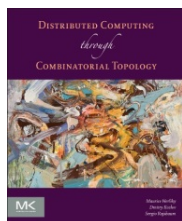
Distributed Computing through
Combinatorial Topology



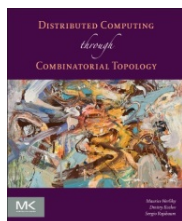
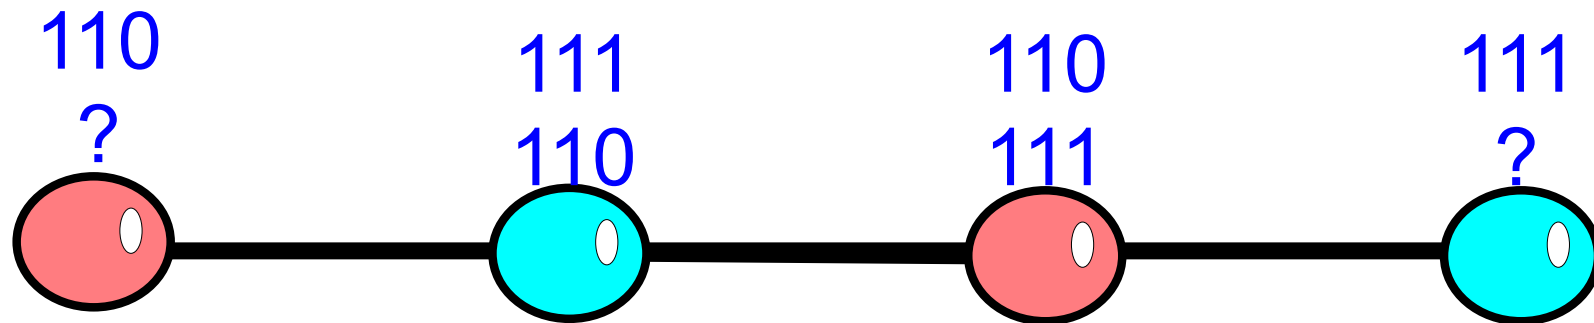


Informally

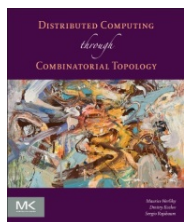
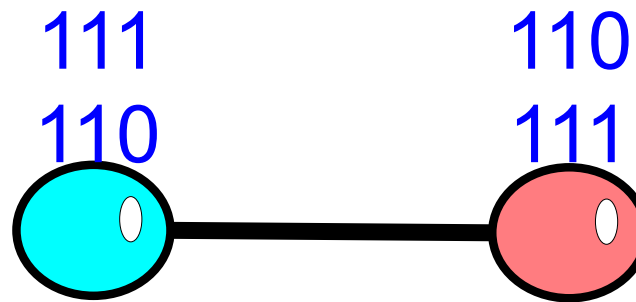
Unreliable communication
does not change “topology” of
global states

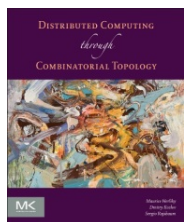
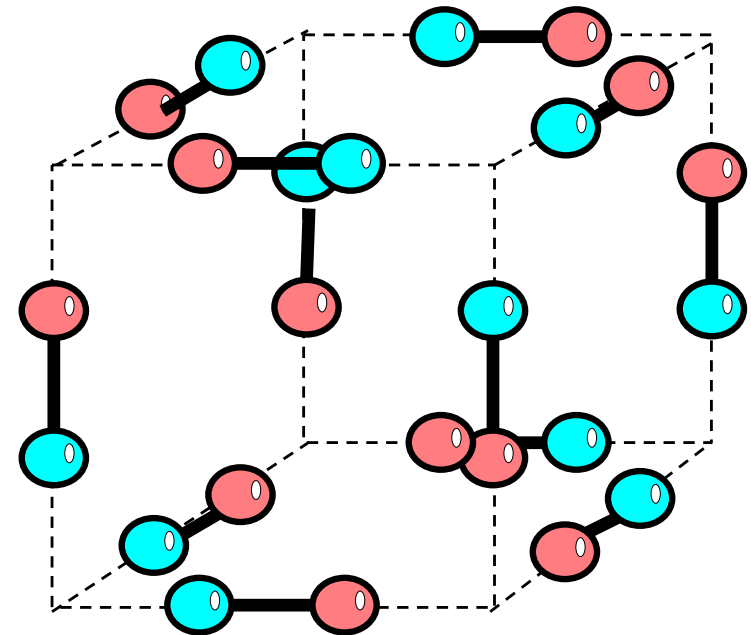
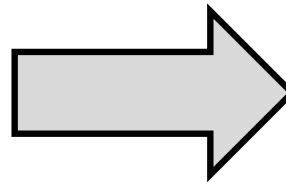
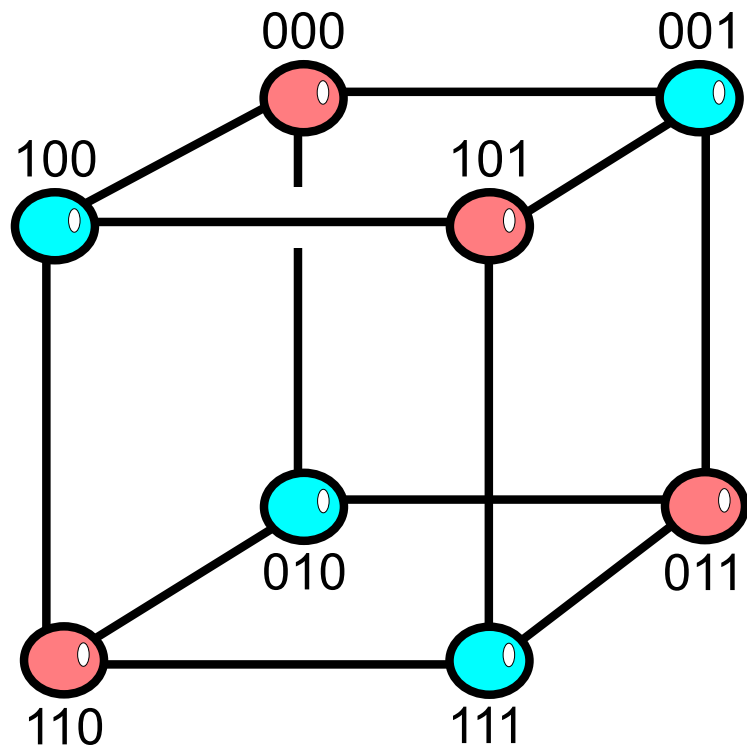


Reliable Communication?



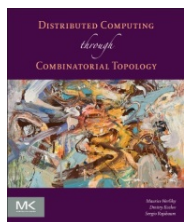
Reliable Communication?





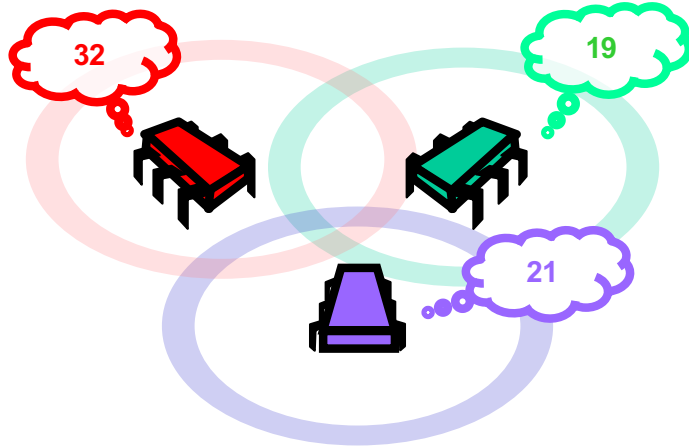
Distributed Computing through
Combinatorial Topology

Tasks

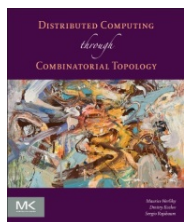


Distributed Computing through
Combinatorial Topology

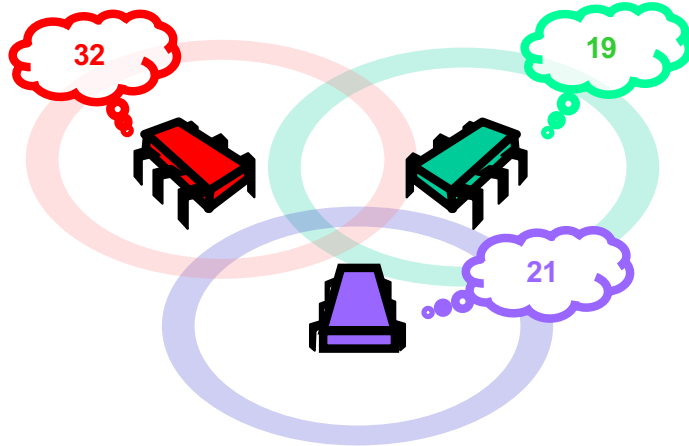
Tasks



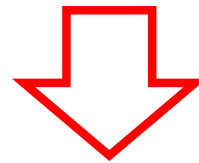
Possible set of *input* values



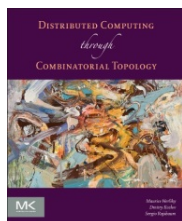
Tasks



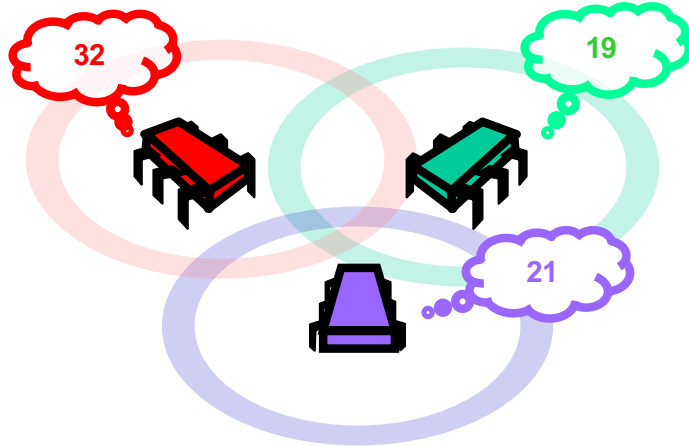
Possible set of *input* values



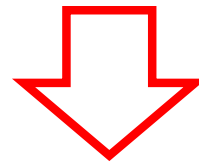
Finite computation



Tasks

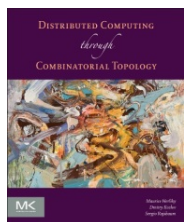
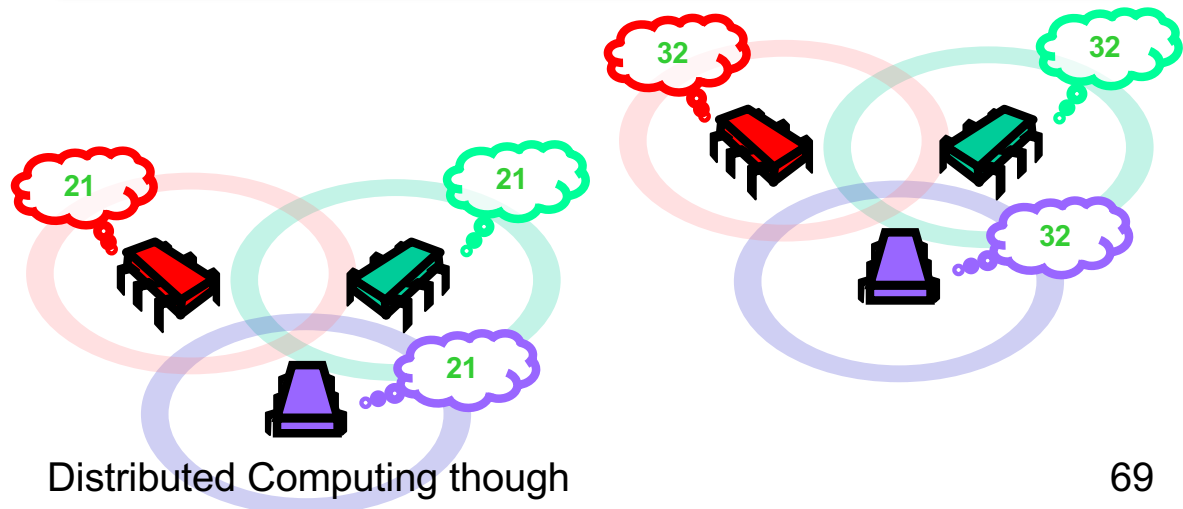
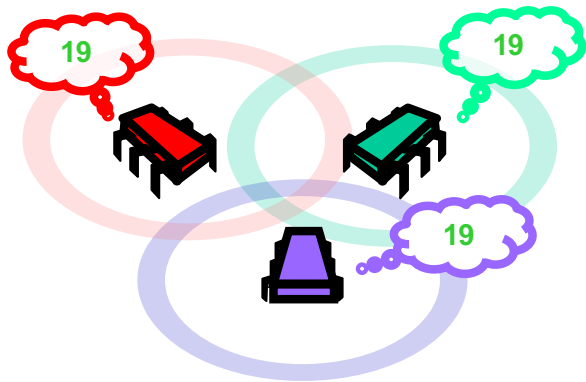


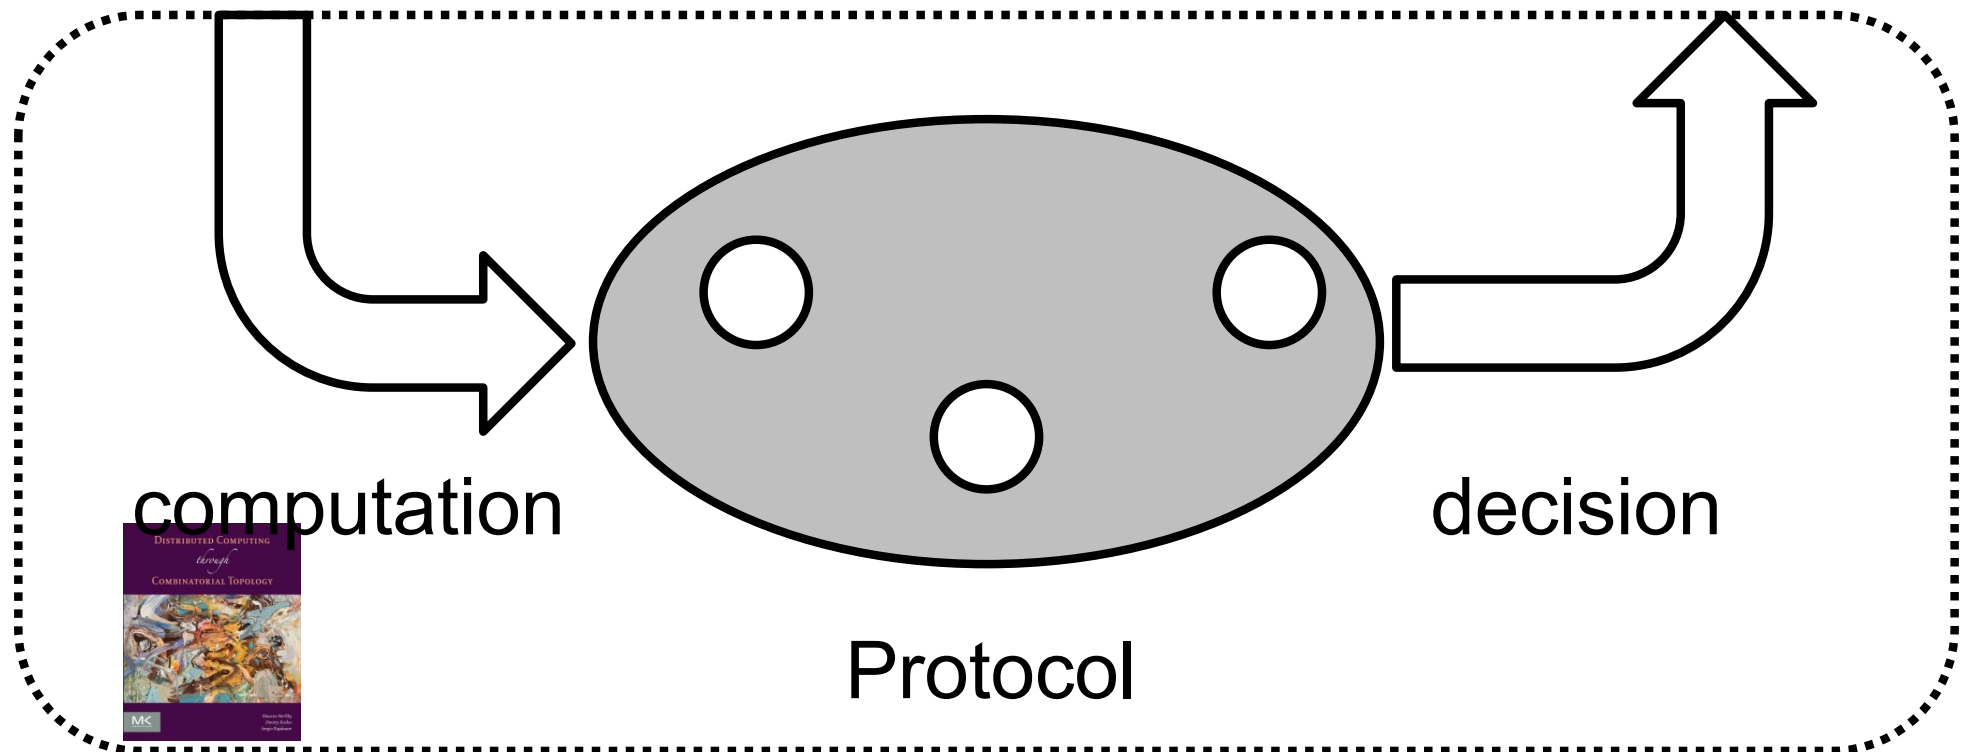
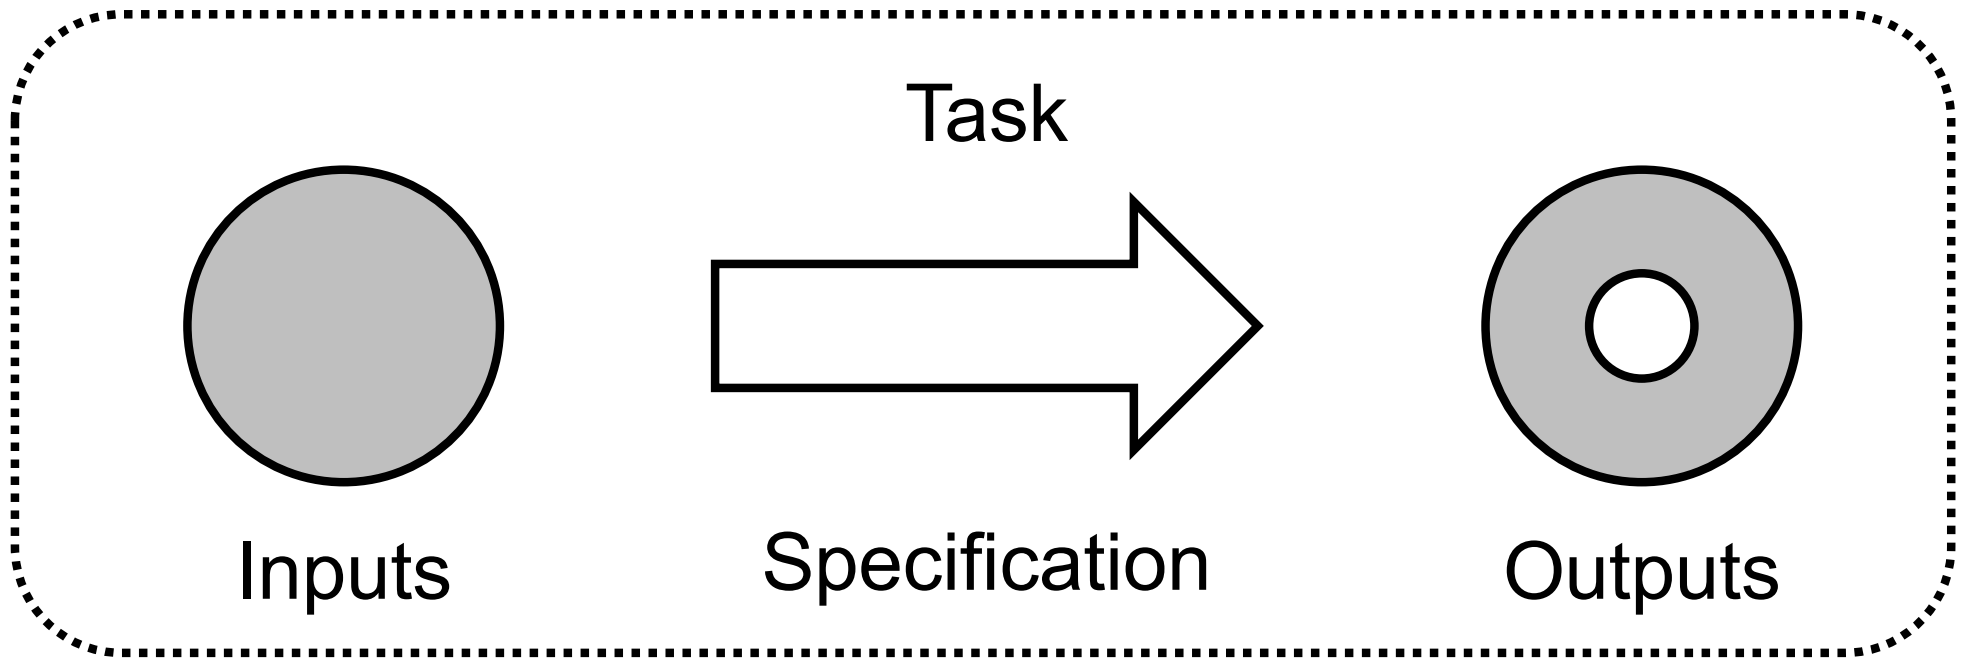
Possible set of *input* values



Finite computation

Possible set of *output* values





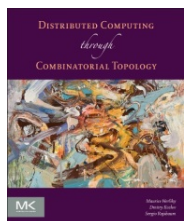
Road Map

Distributed Computing

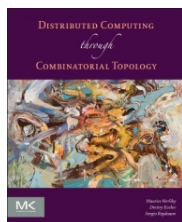
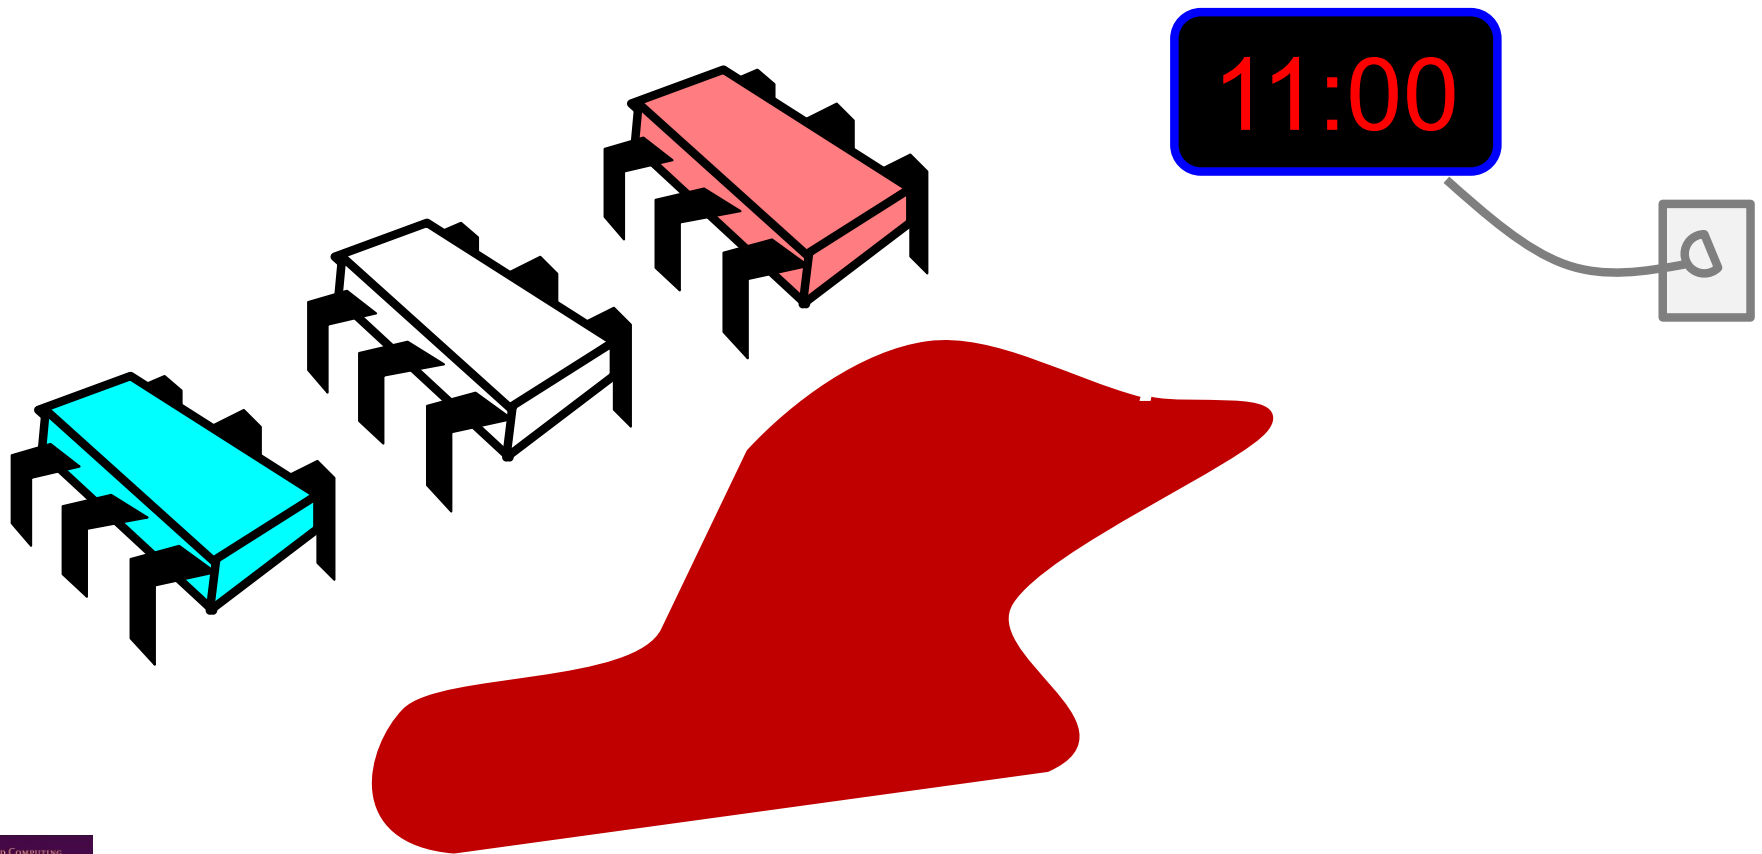
Two Classic Distributed Problems

The Muddy Children

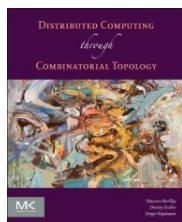
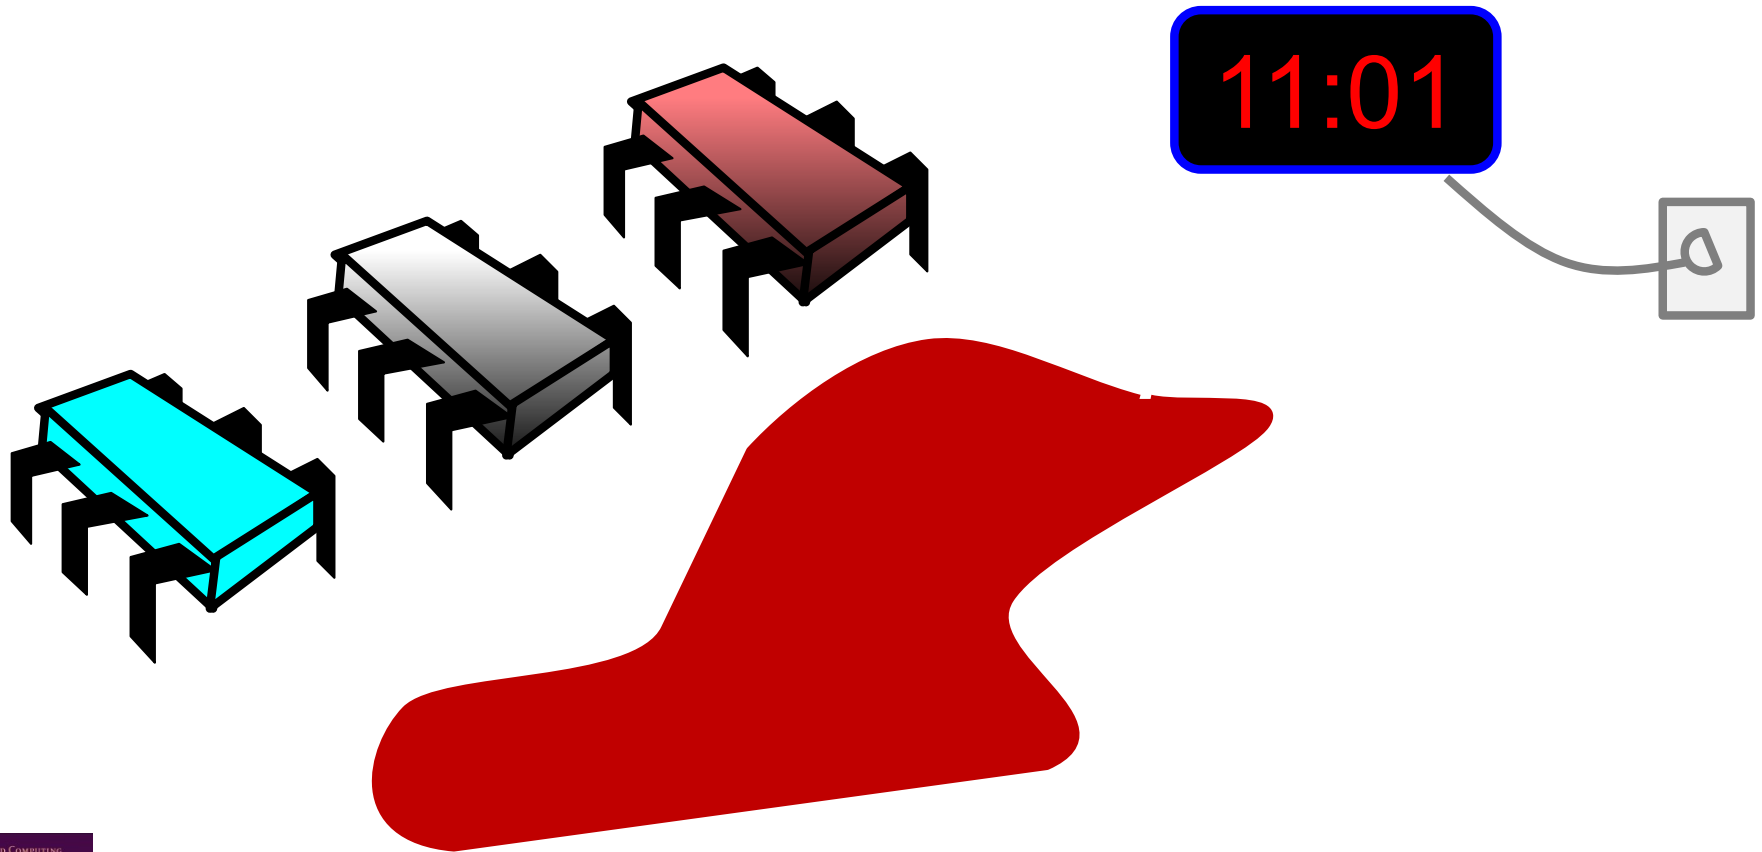
Coordinated Attack



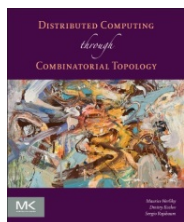
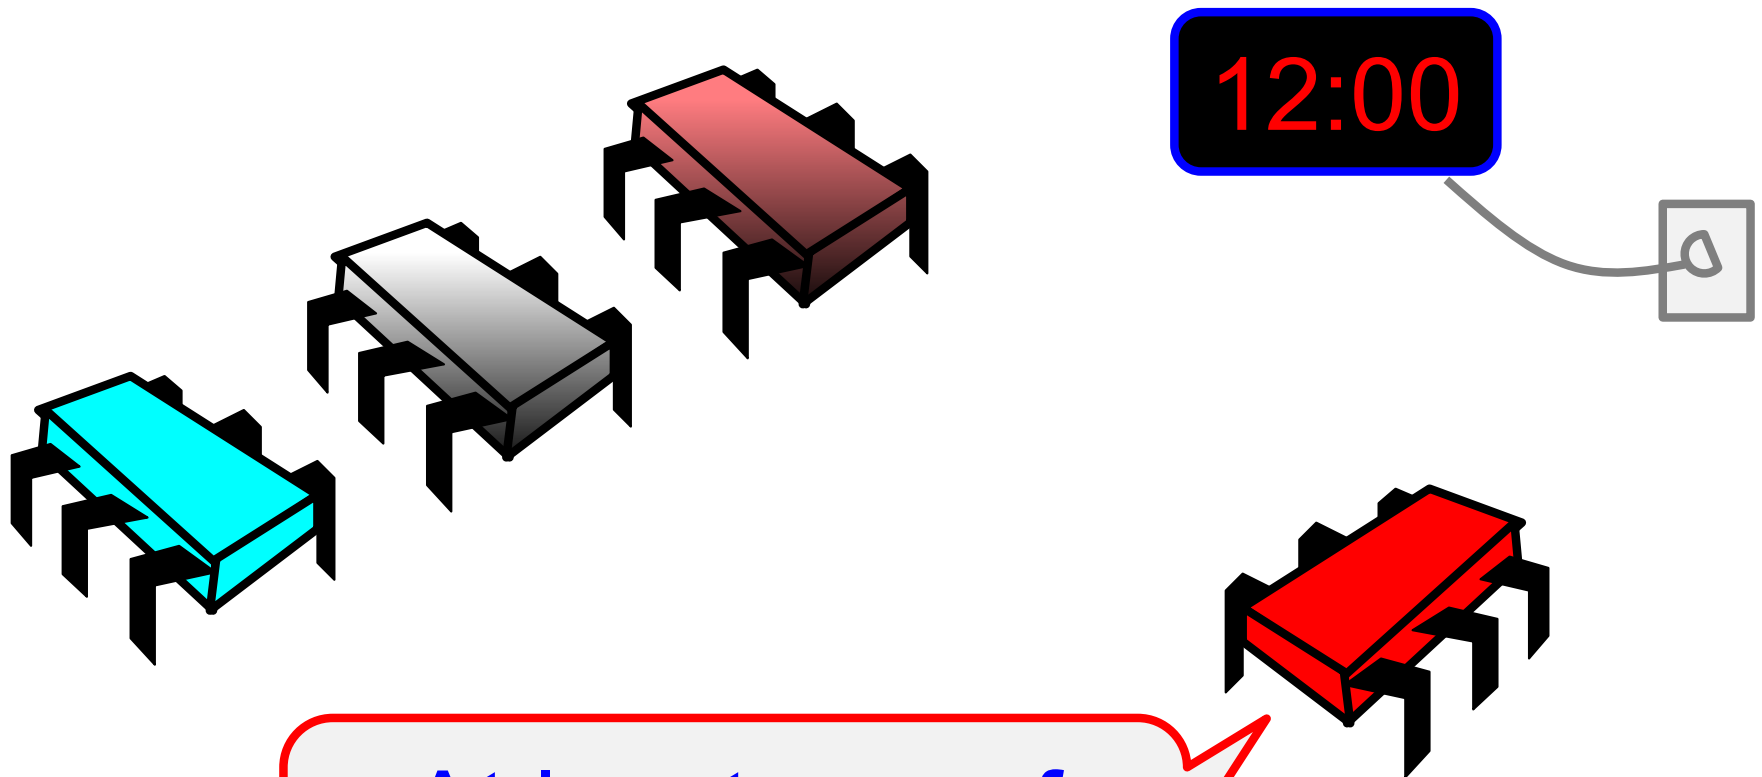
Muddy Children



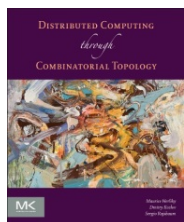
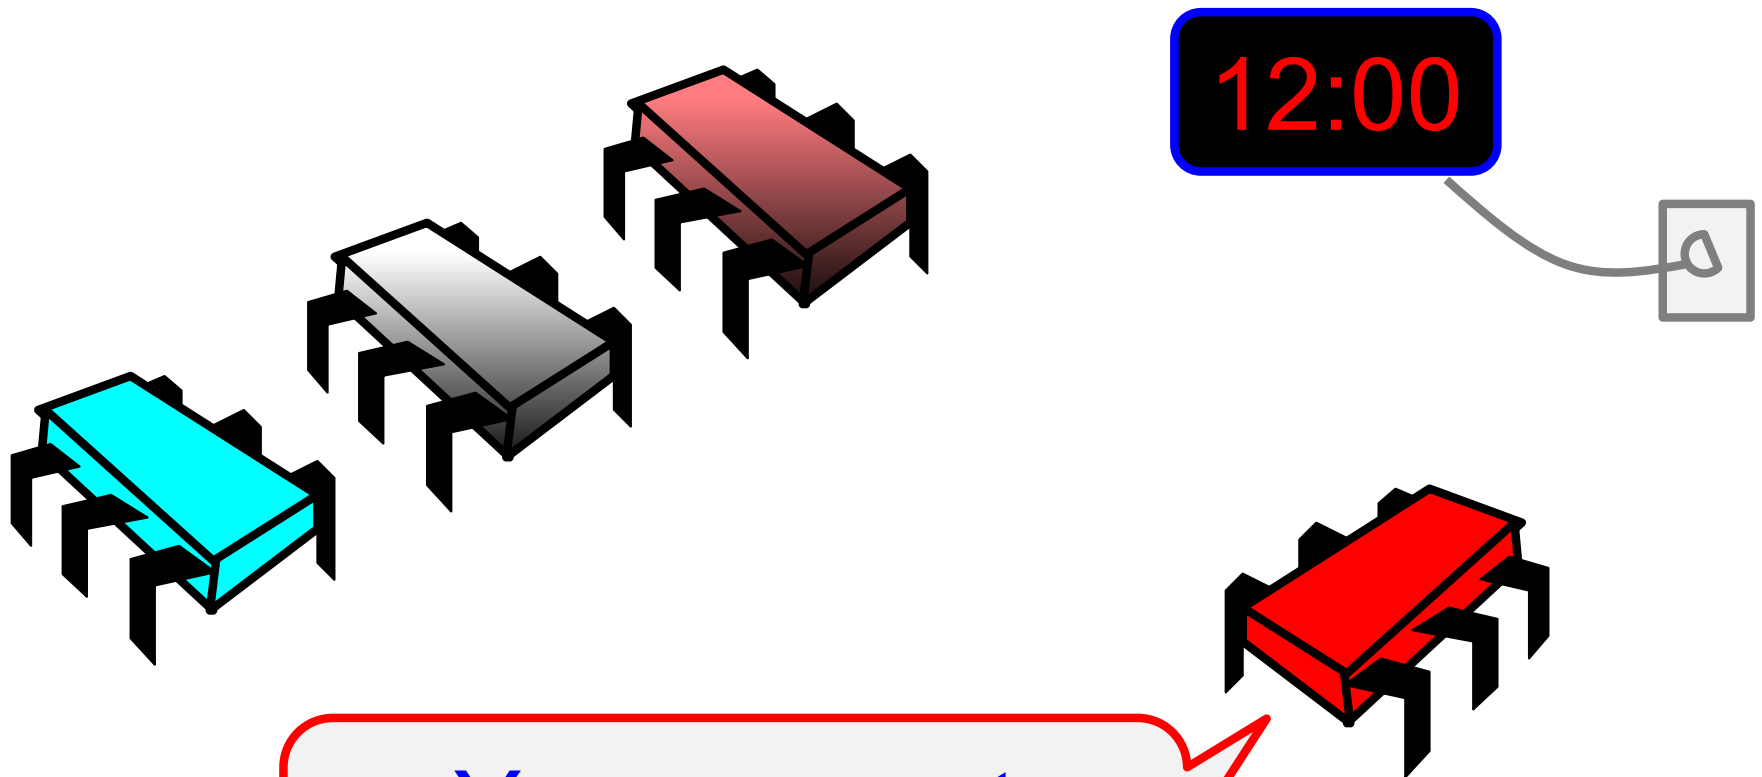
Muddy Children



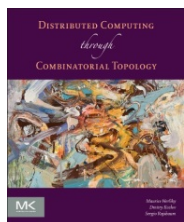
Muddy Children



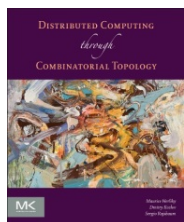
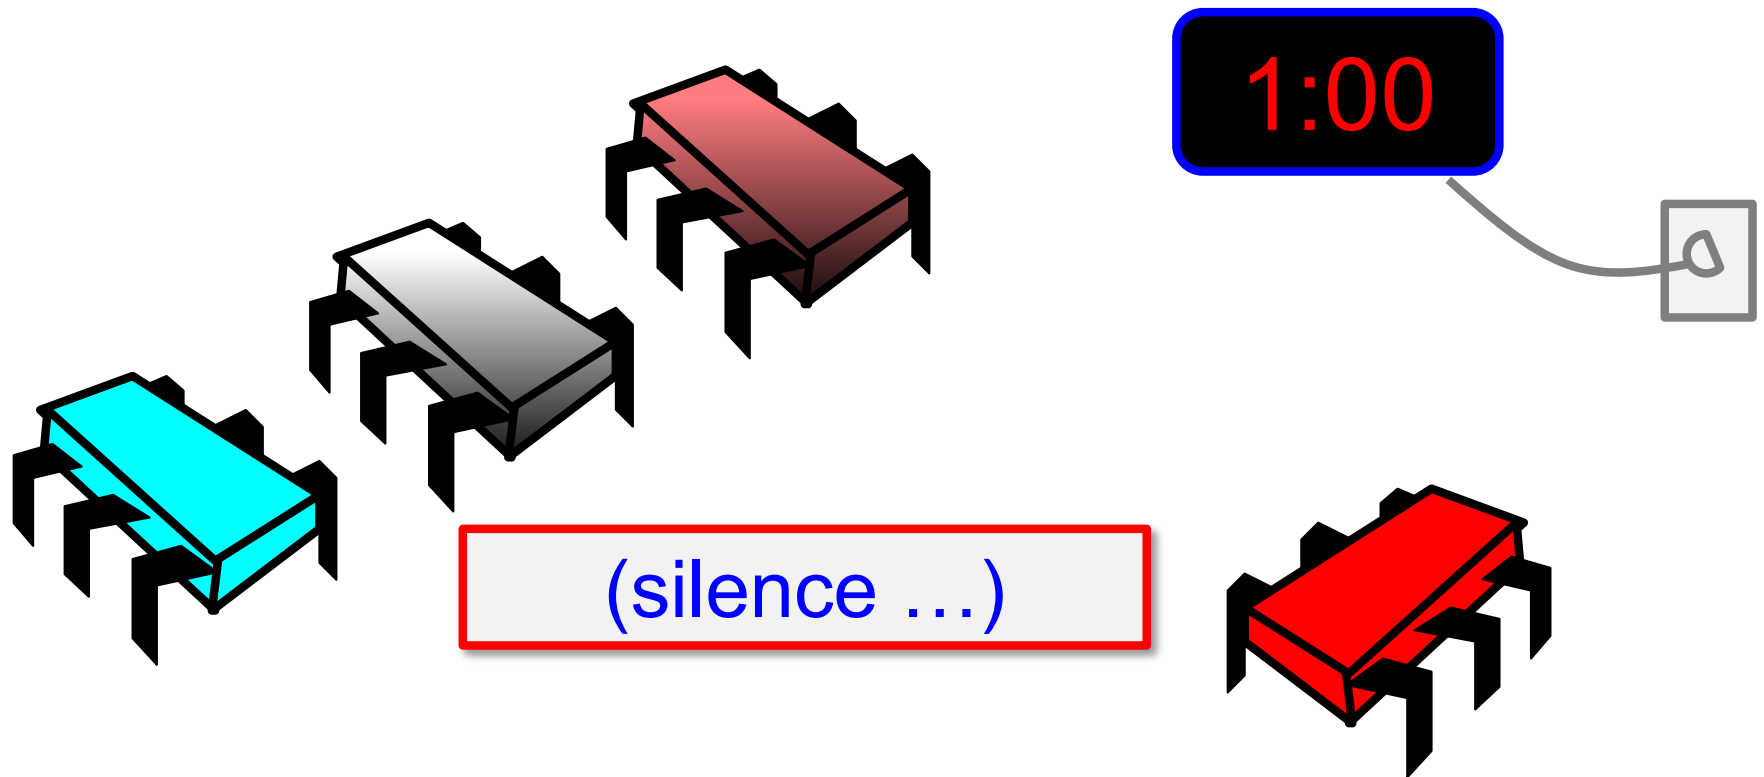
Muddy Children



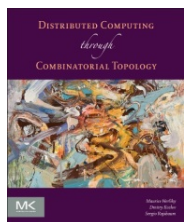
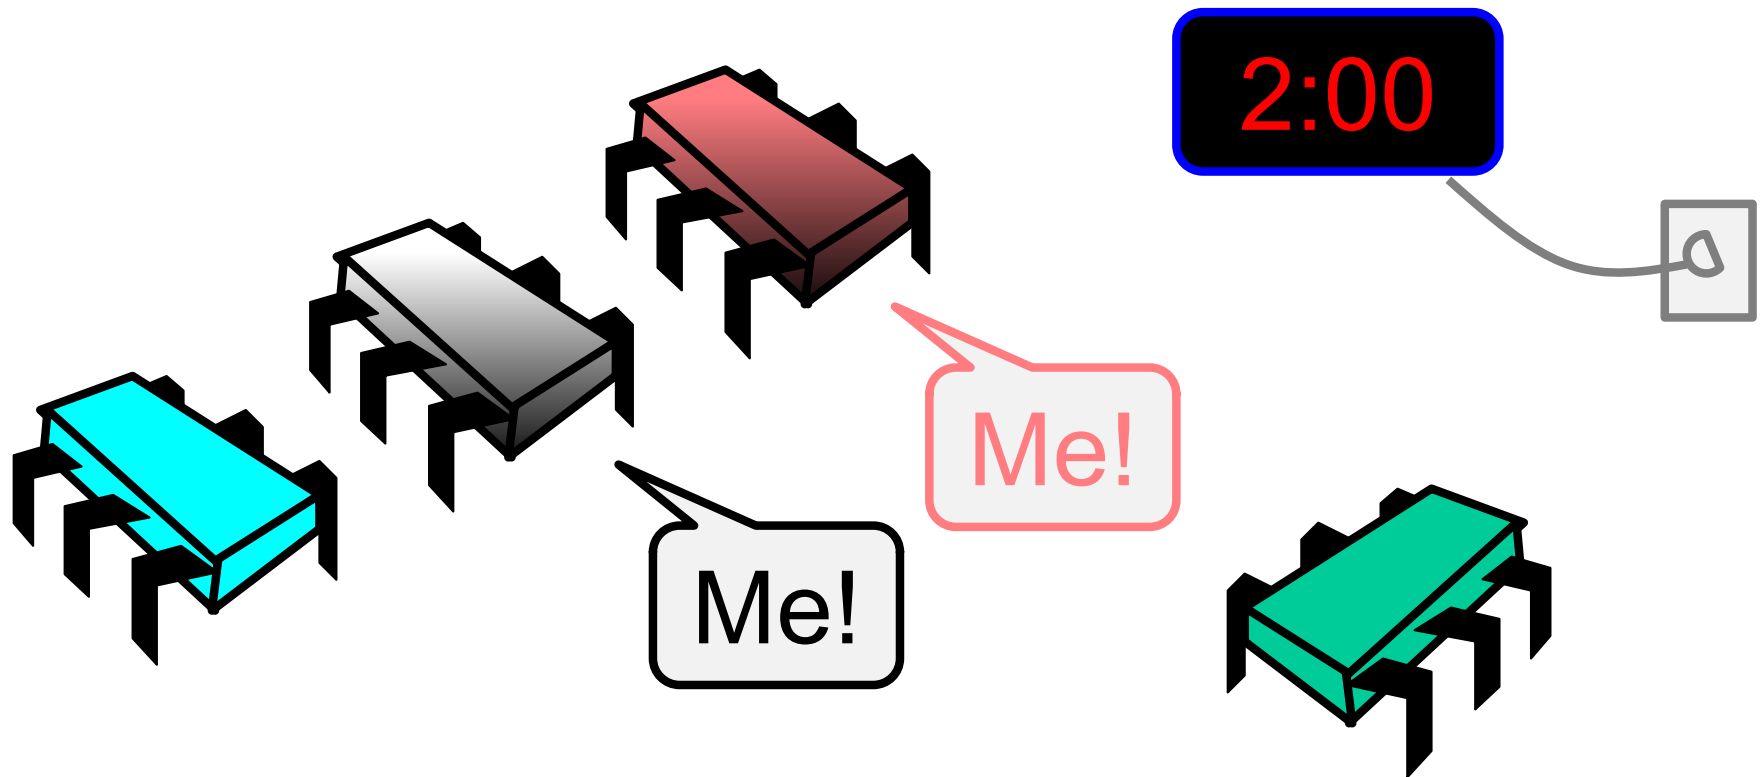
Muddy Children



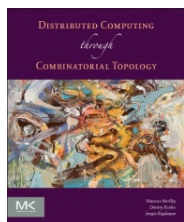
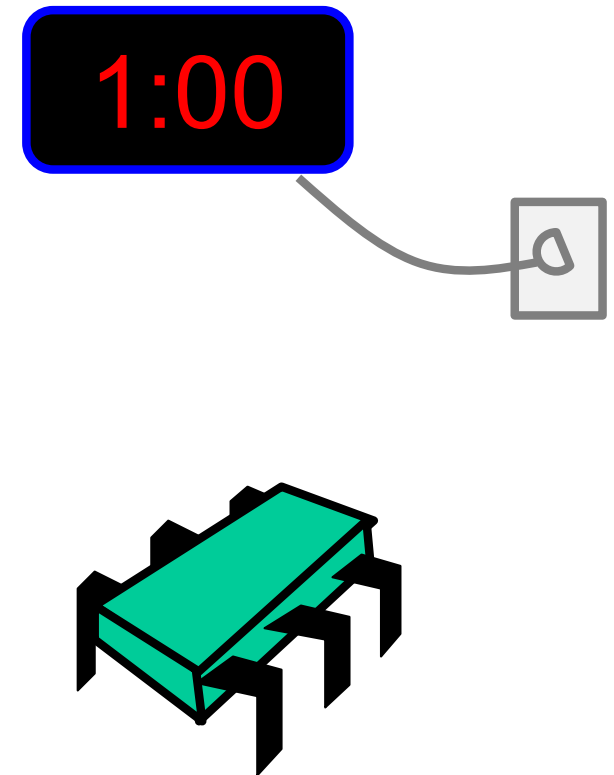
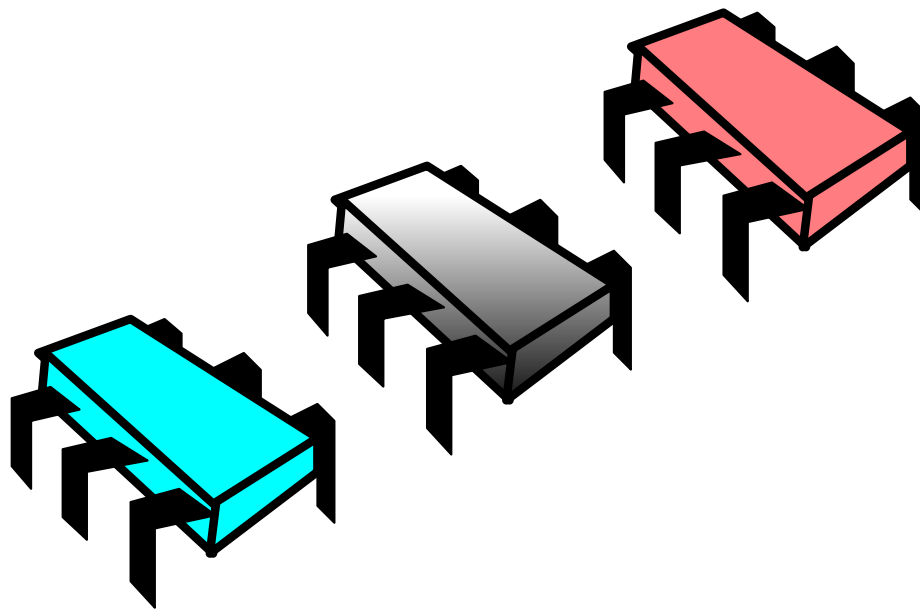
Muddy Children

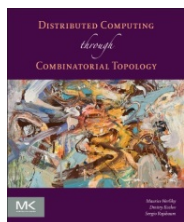


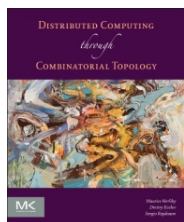
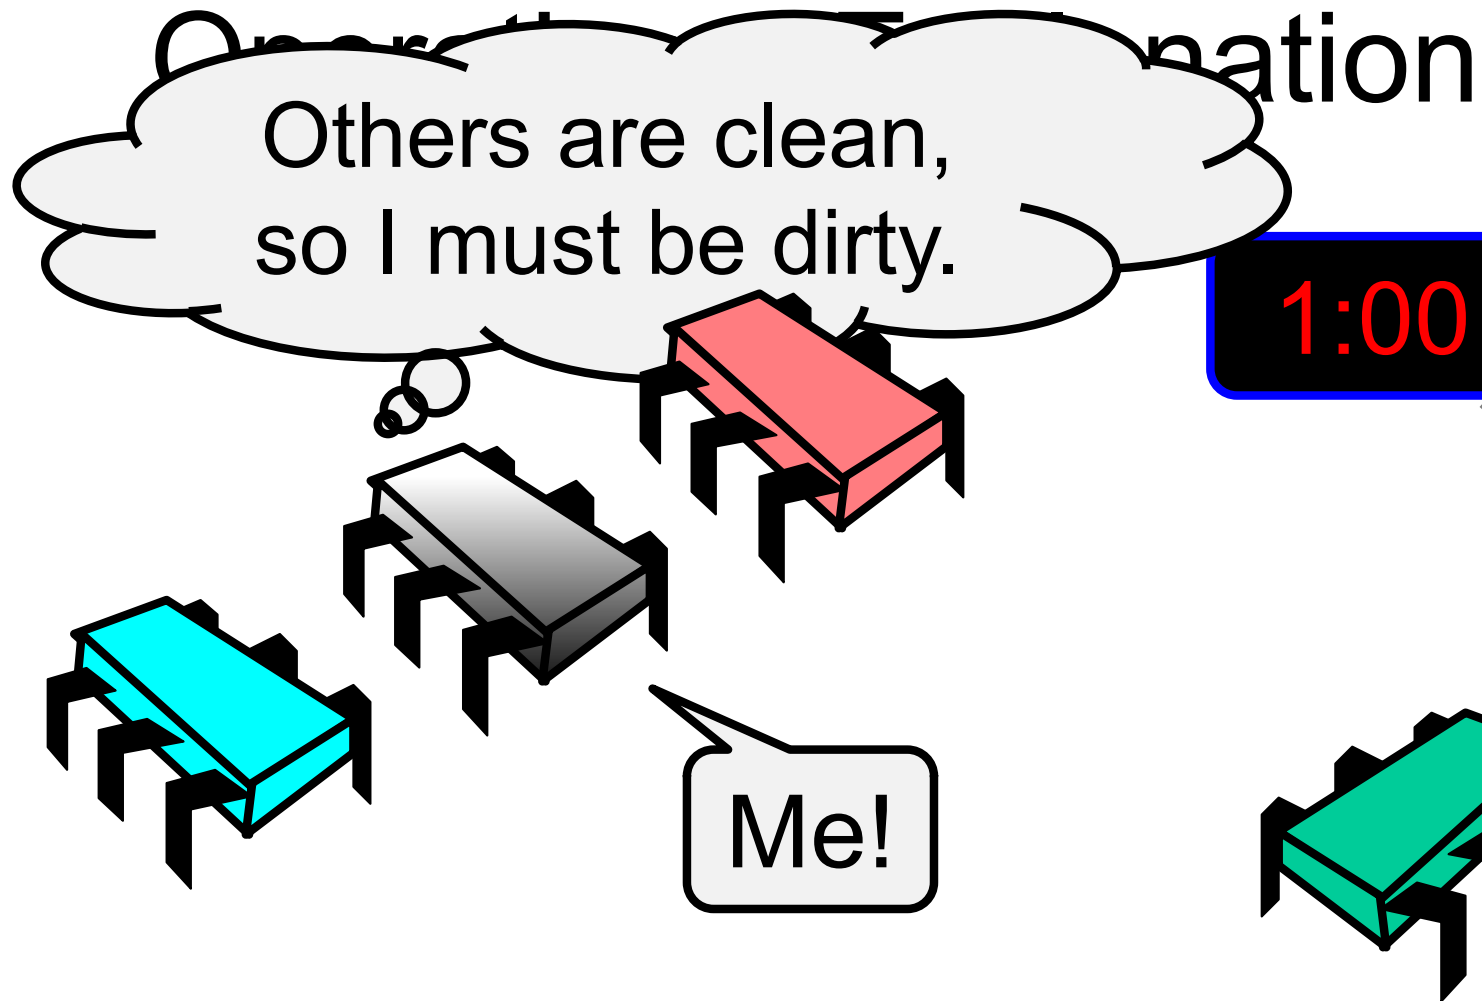
Muddy Children



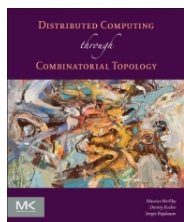
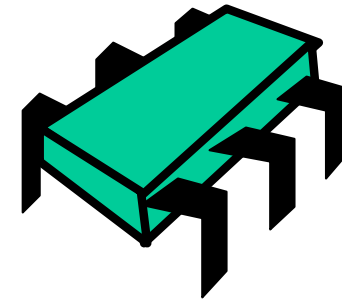
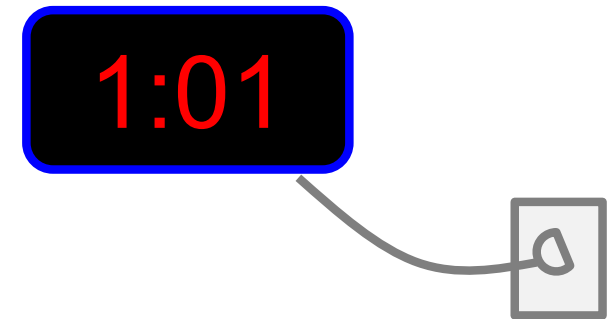
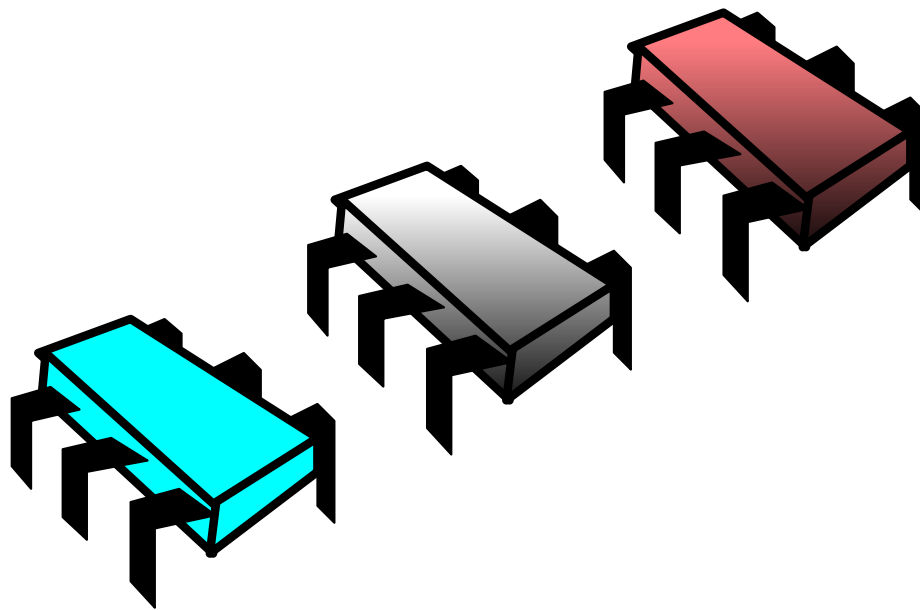
Operational Explanation







Operational Explanation

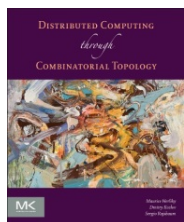


Operational Explanation

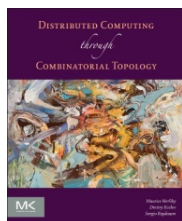
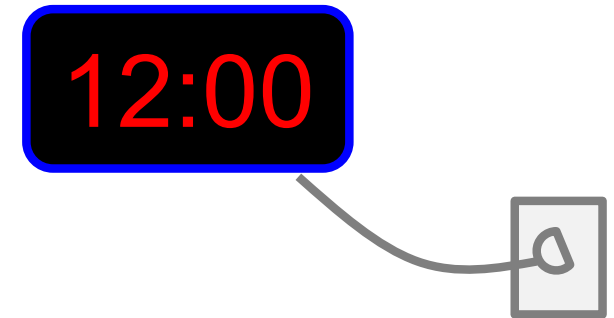
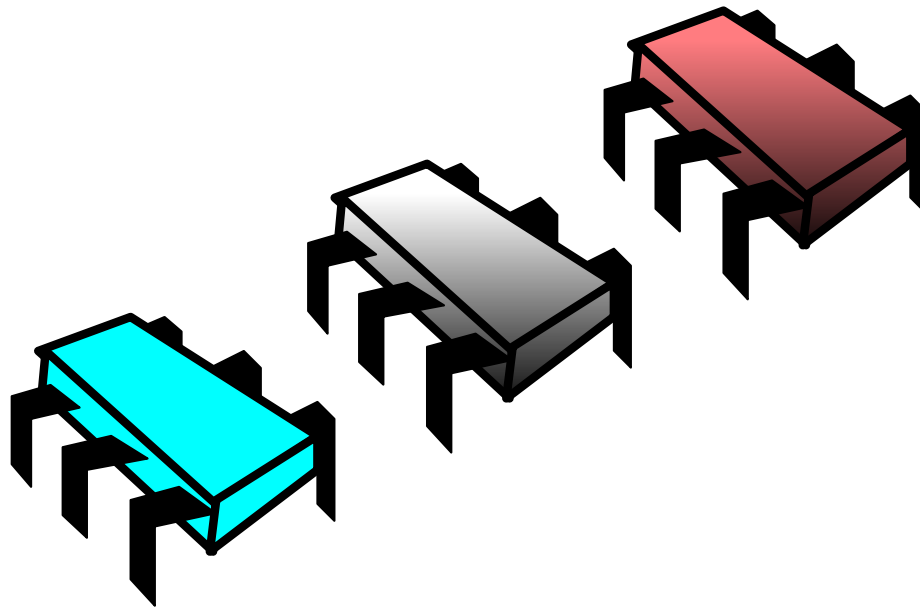
He was quiet, so
I must be dirty.

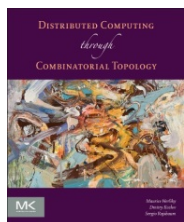
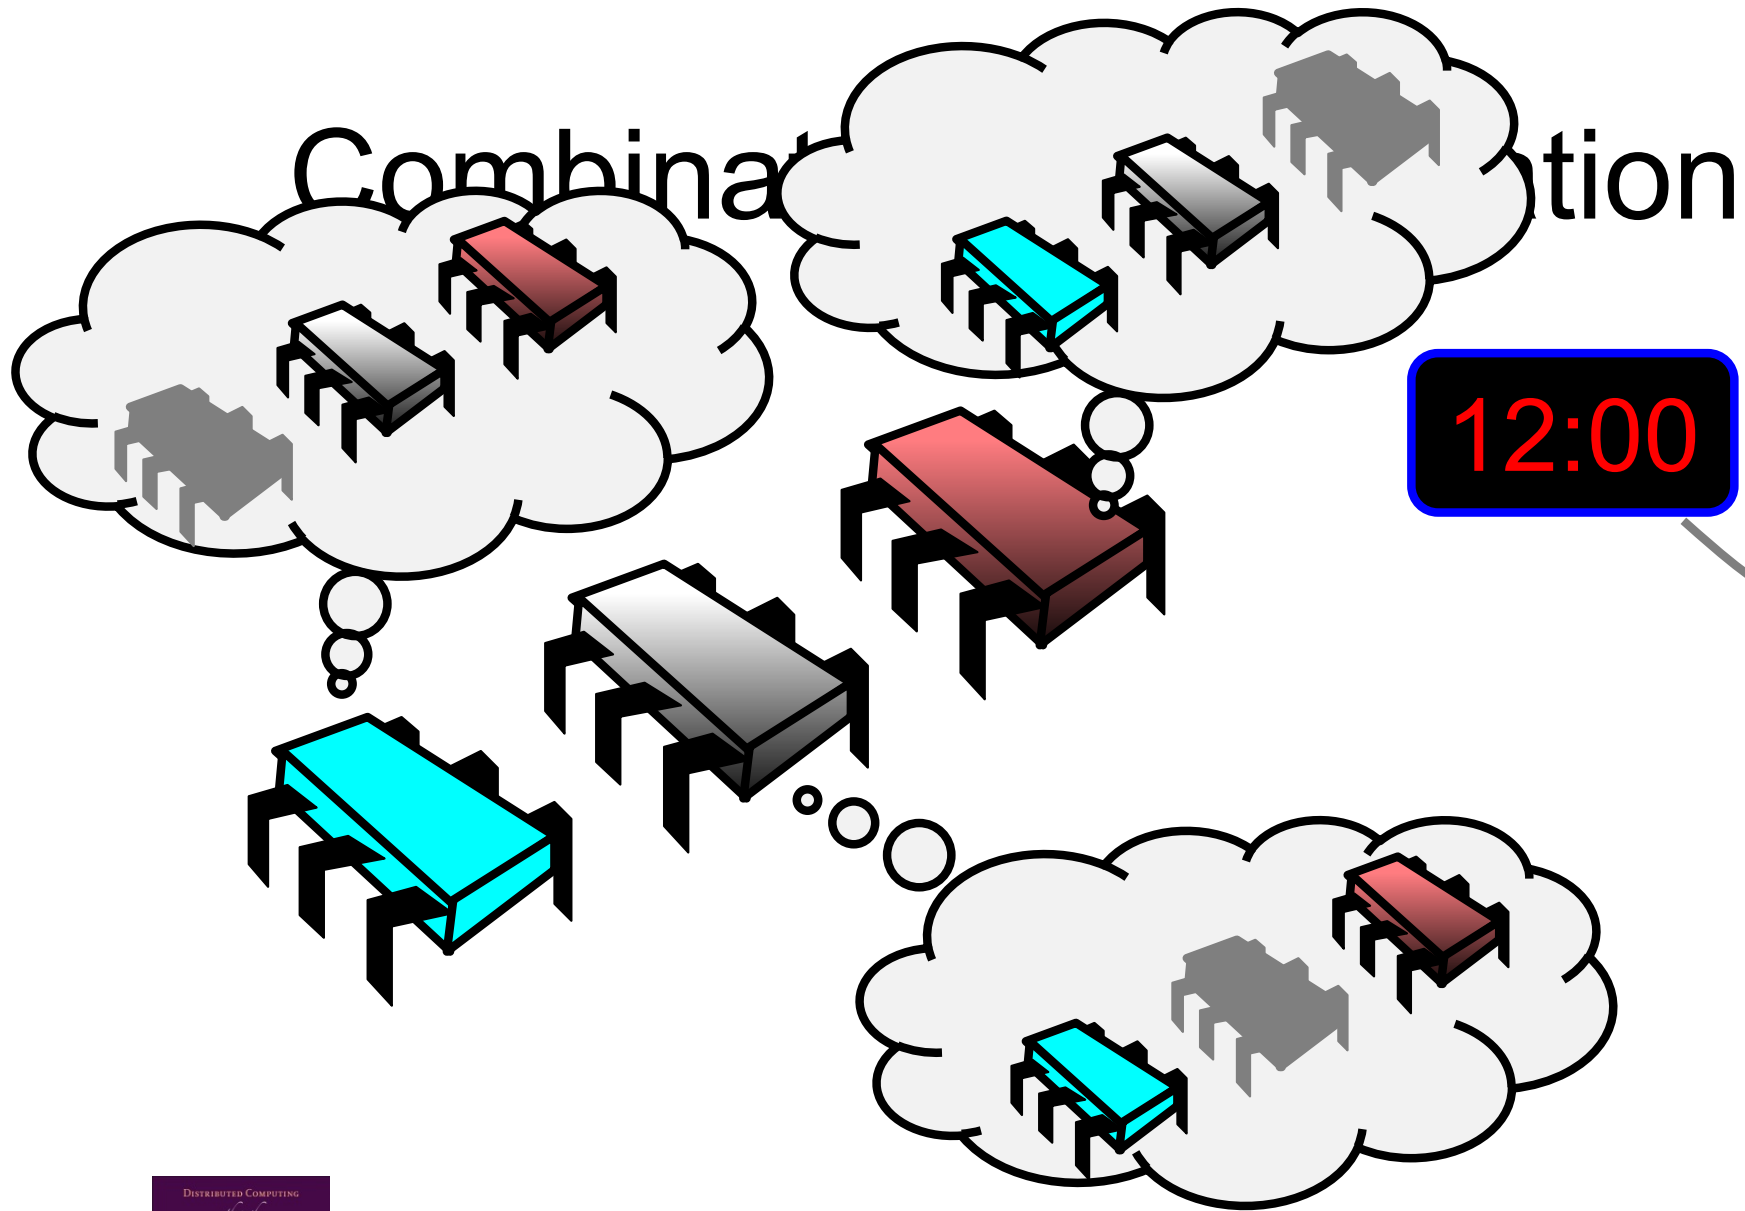
1:01

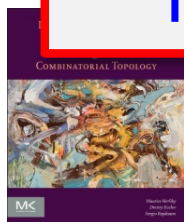
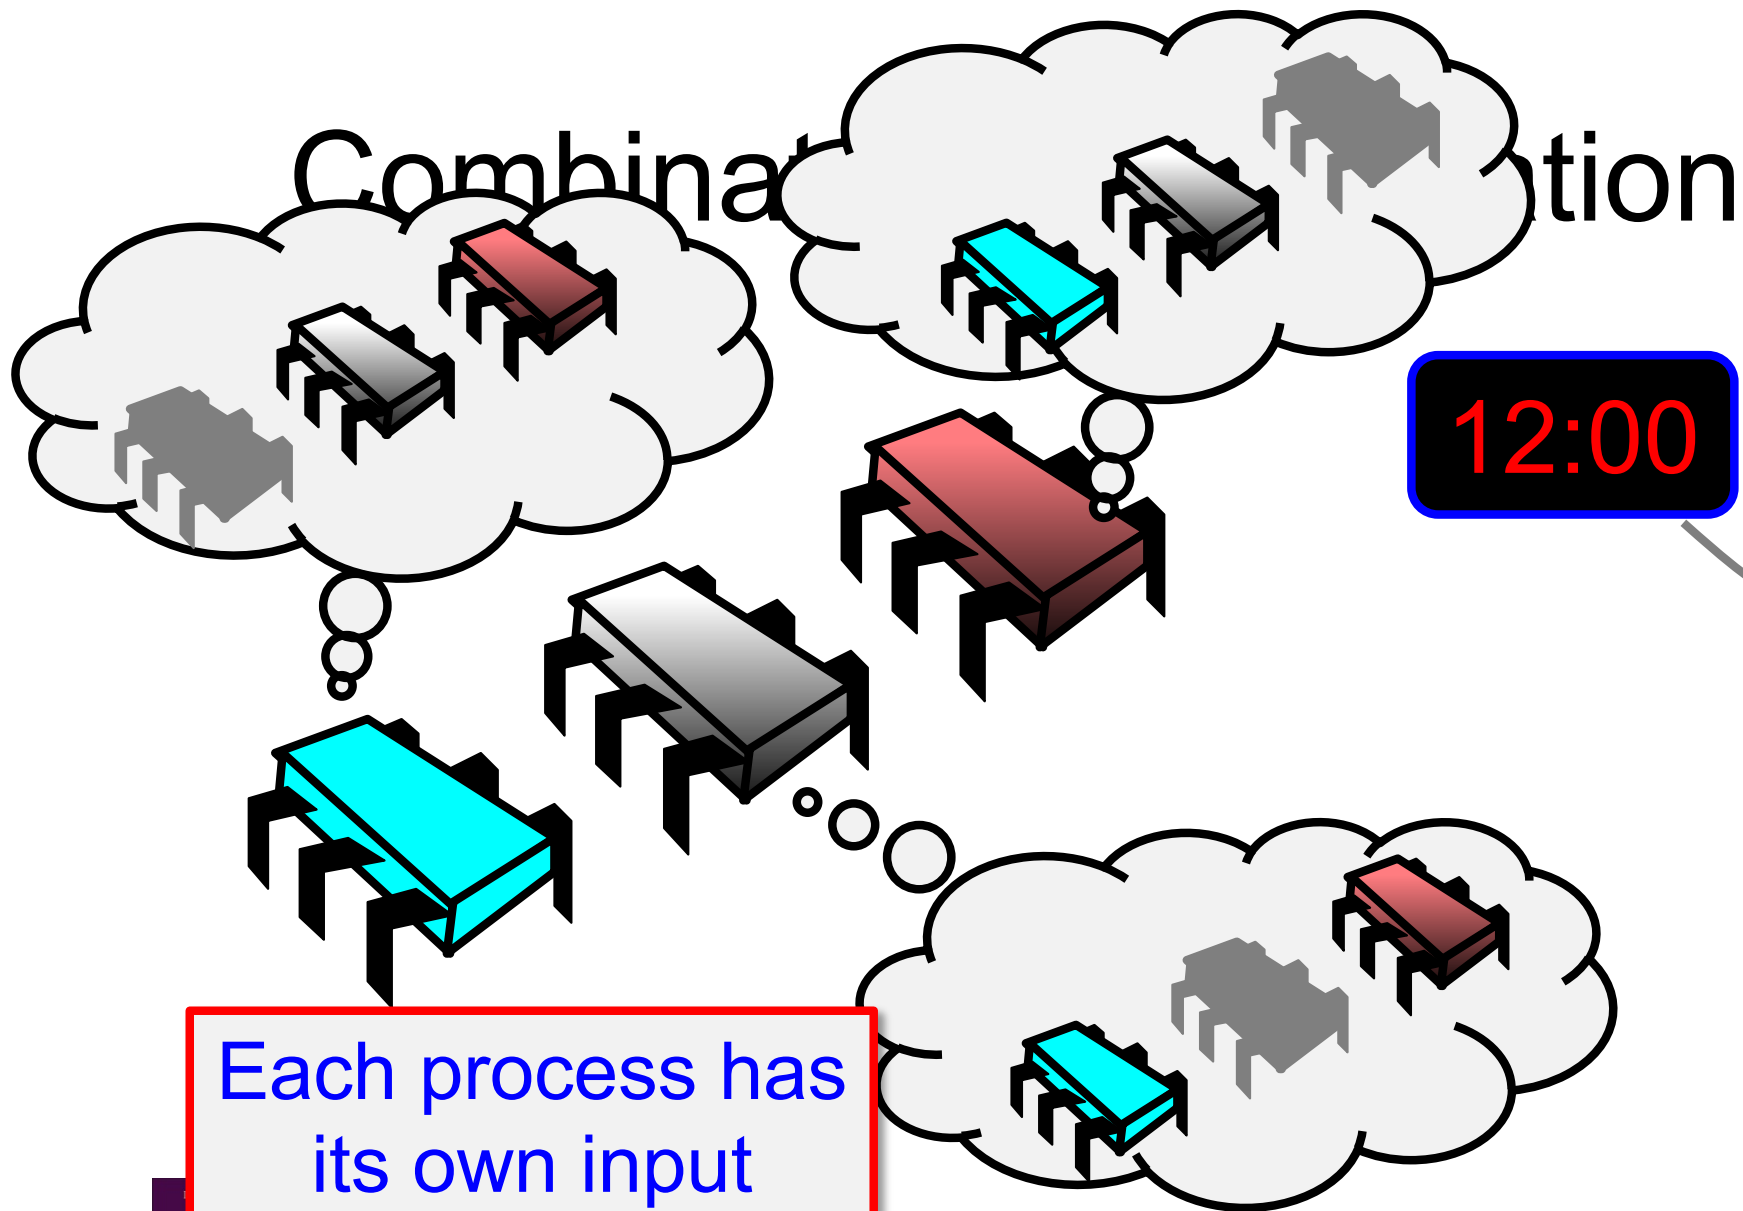
He was quiet, so
I must be dirty.



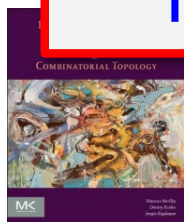
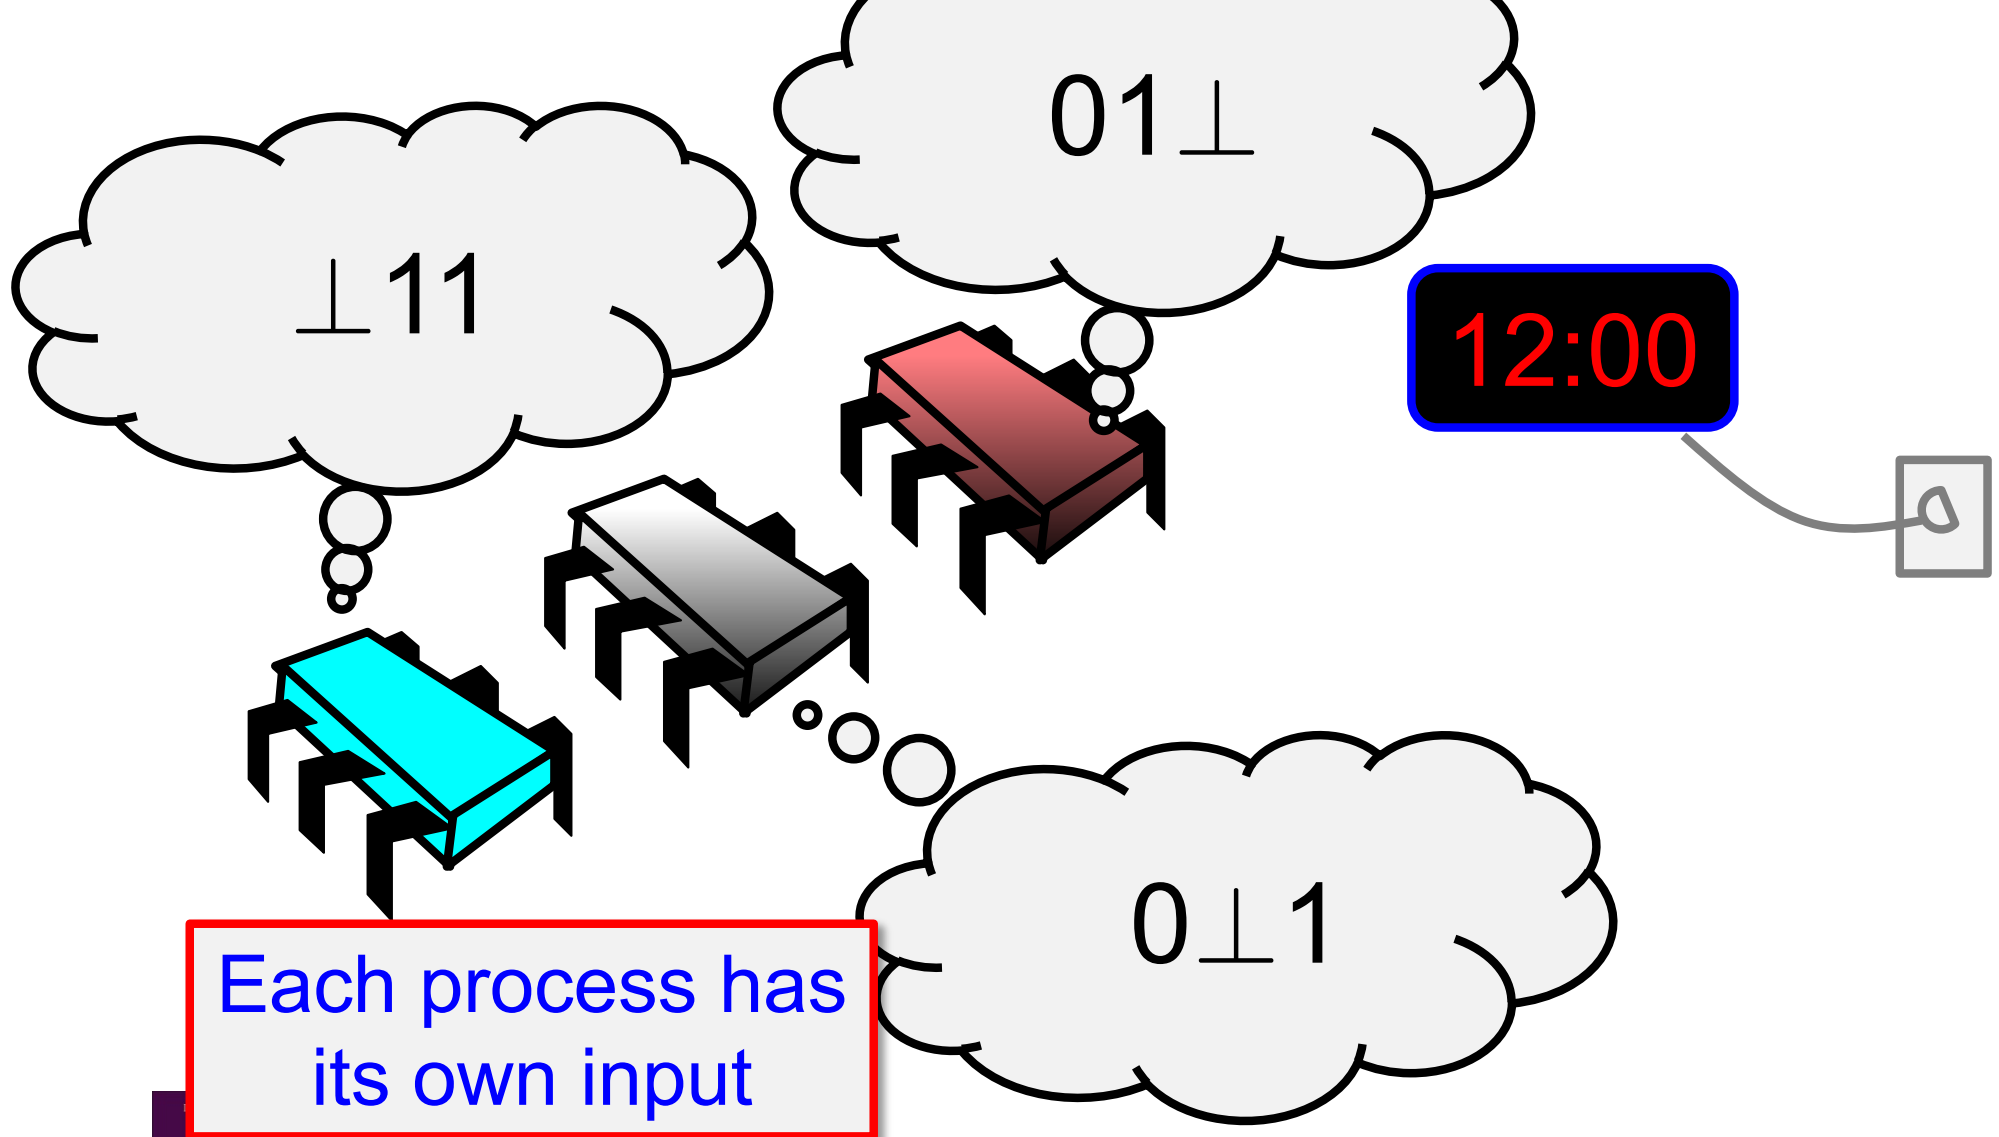
Combinatorial Explanation



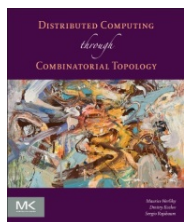
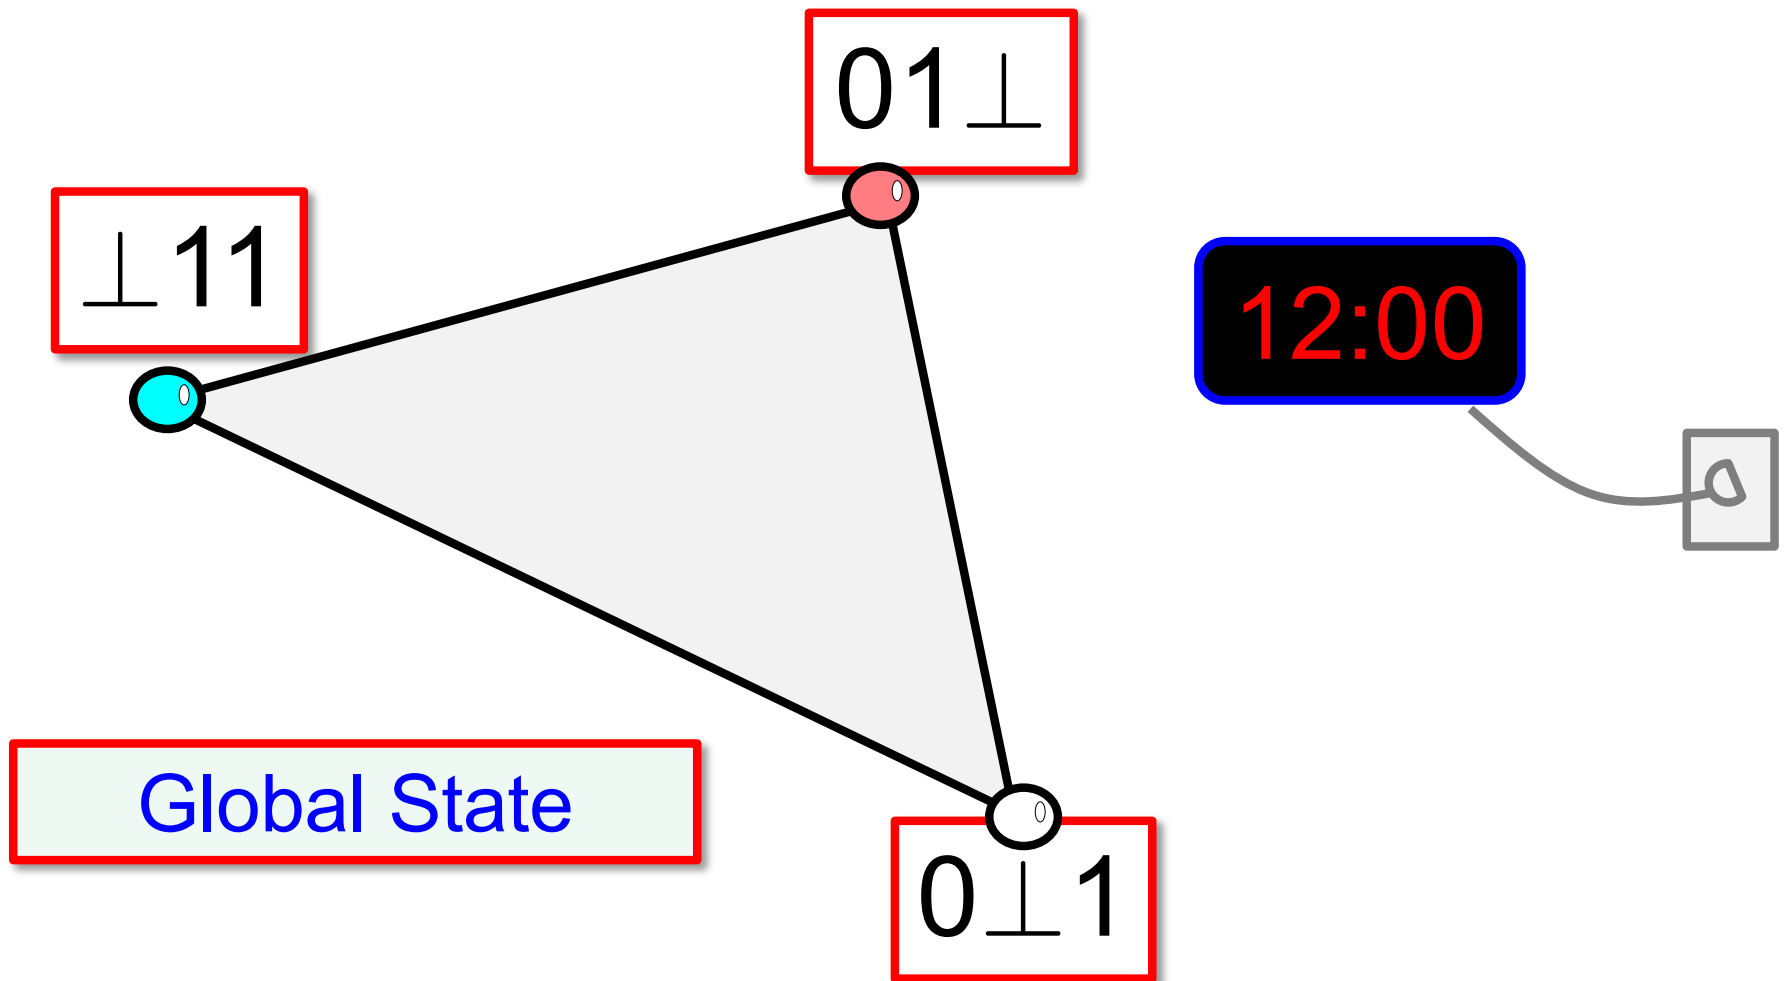


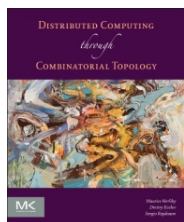
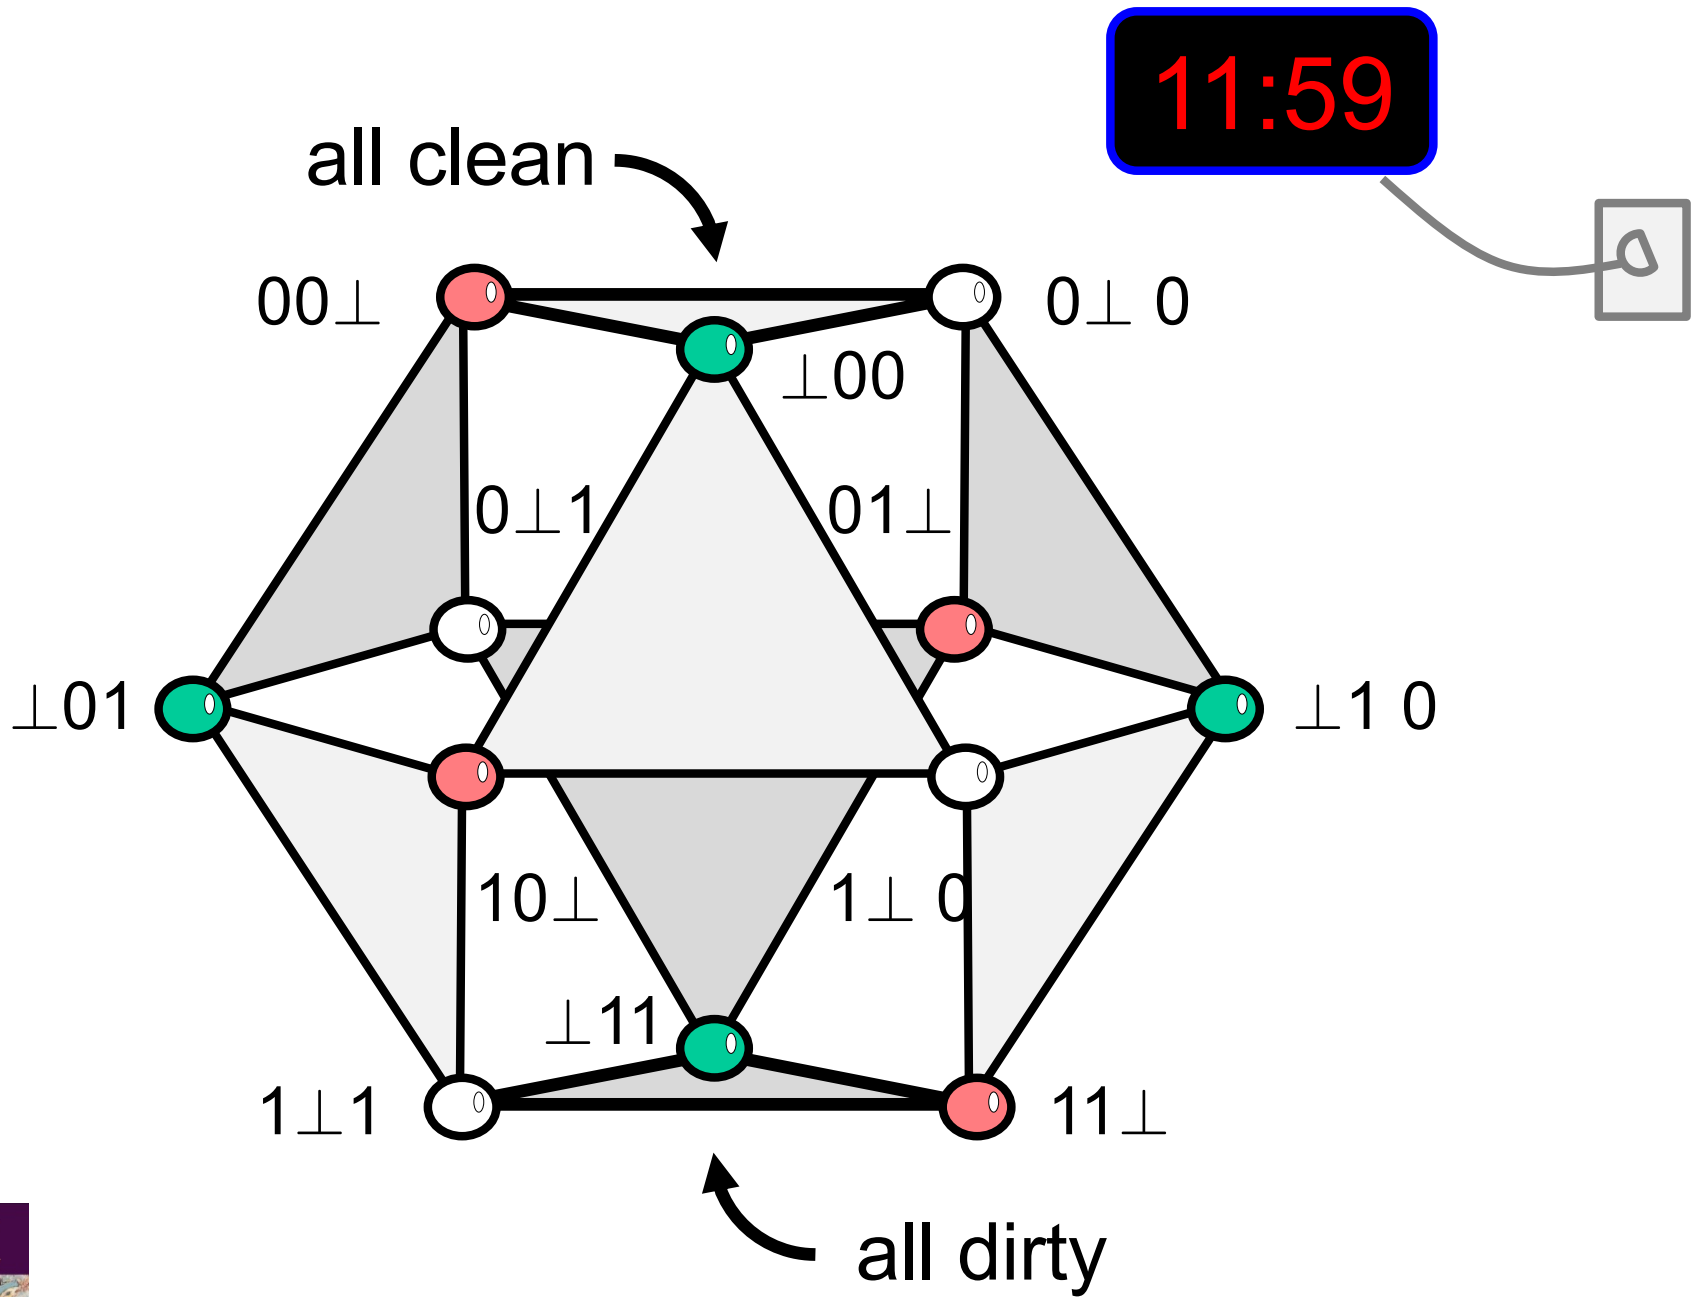


Combinatorial Enumeration



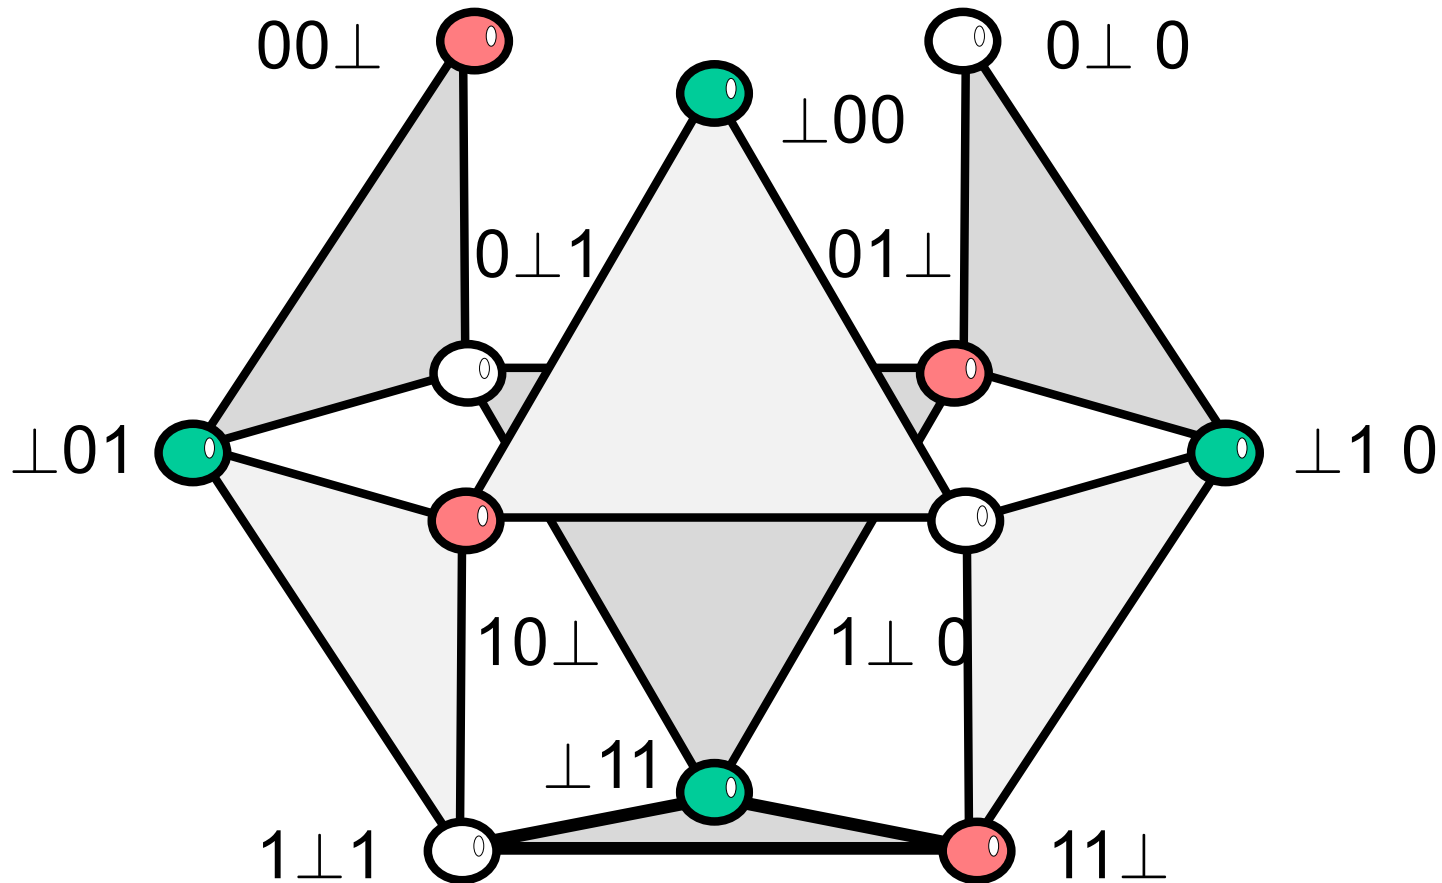
Combinatorial Explanation



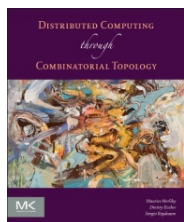


Distributed Computing through
Combinatorial Topology

12:01

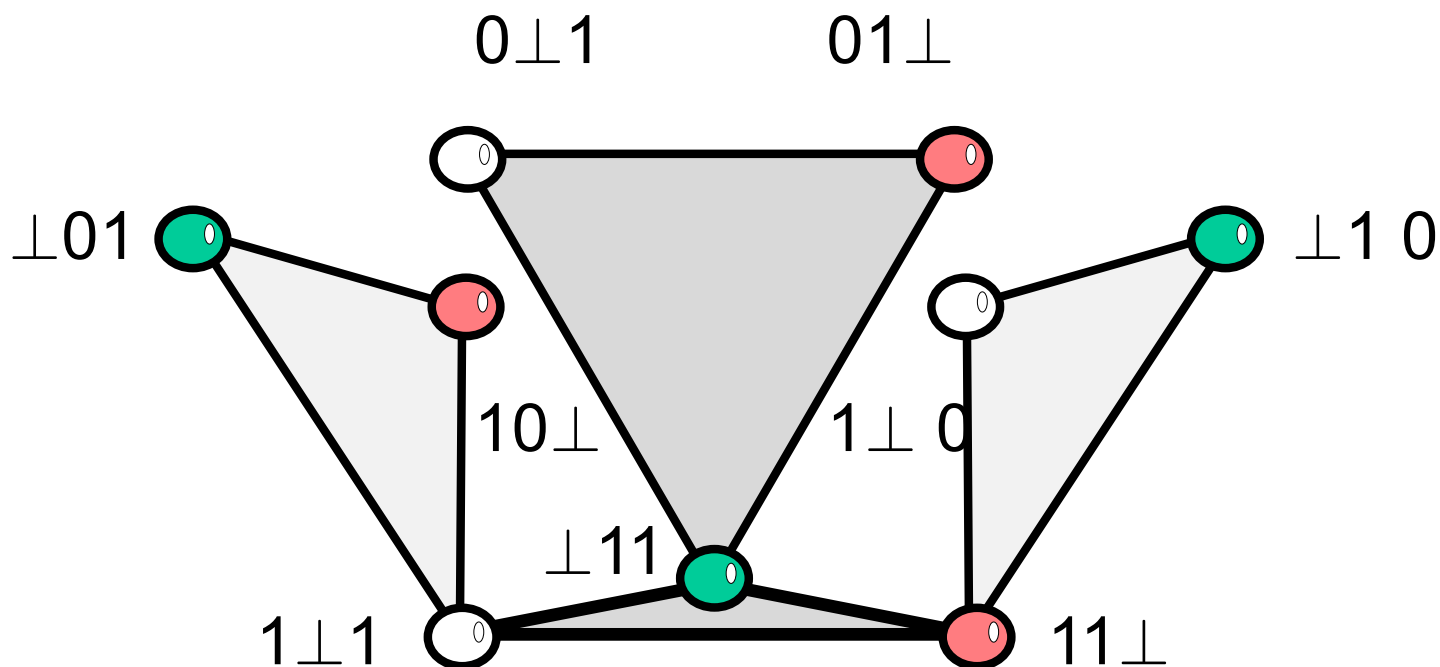


all dirty

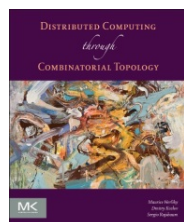


Distributed Computing through
Combinatorial Topology

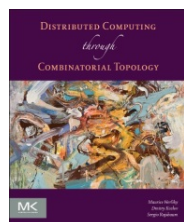
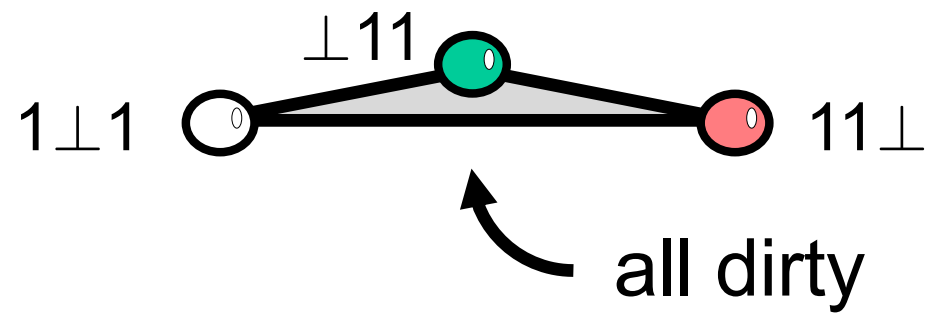
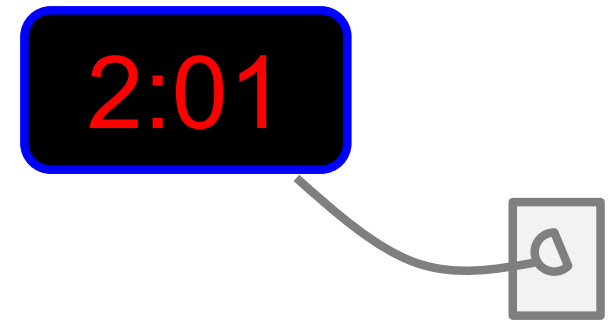
1:01



all dirty



Distributed Computing through
Combinatorial Topology



Distributed Computing through
Combinatorial Topology

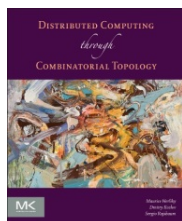
Road Map

Distributed Computing

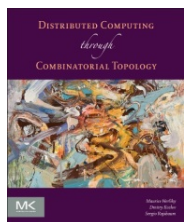
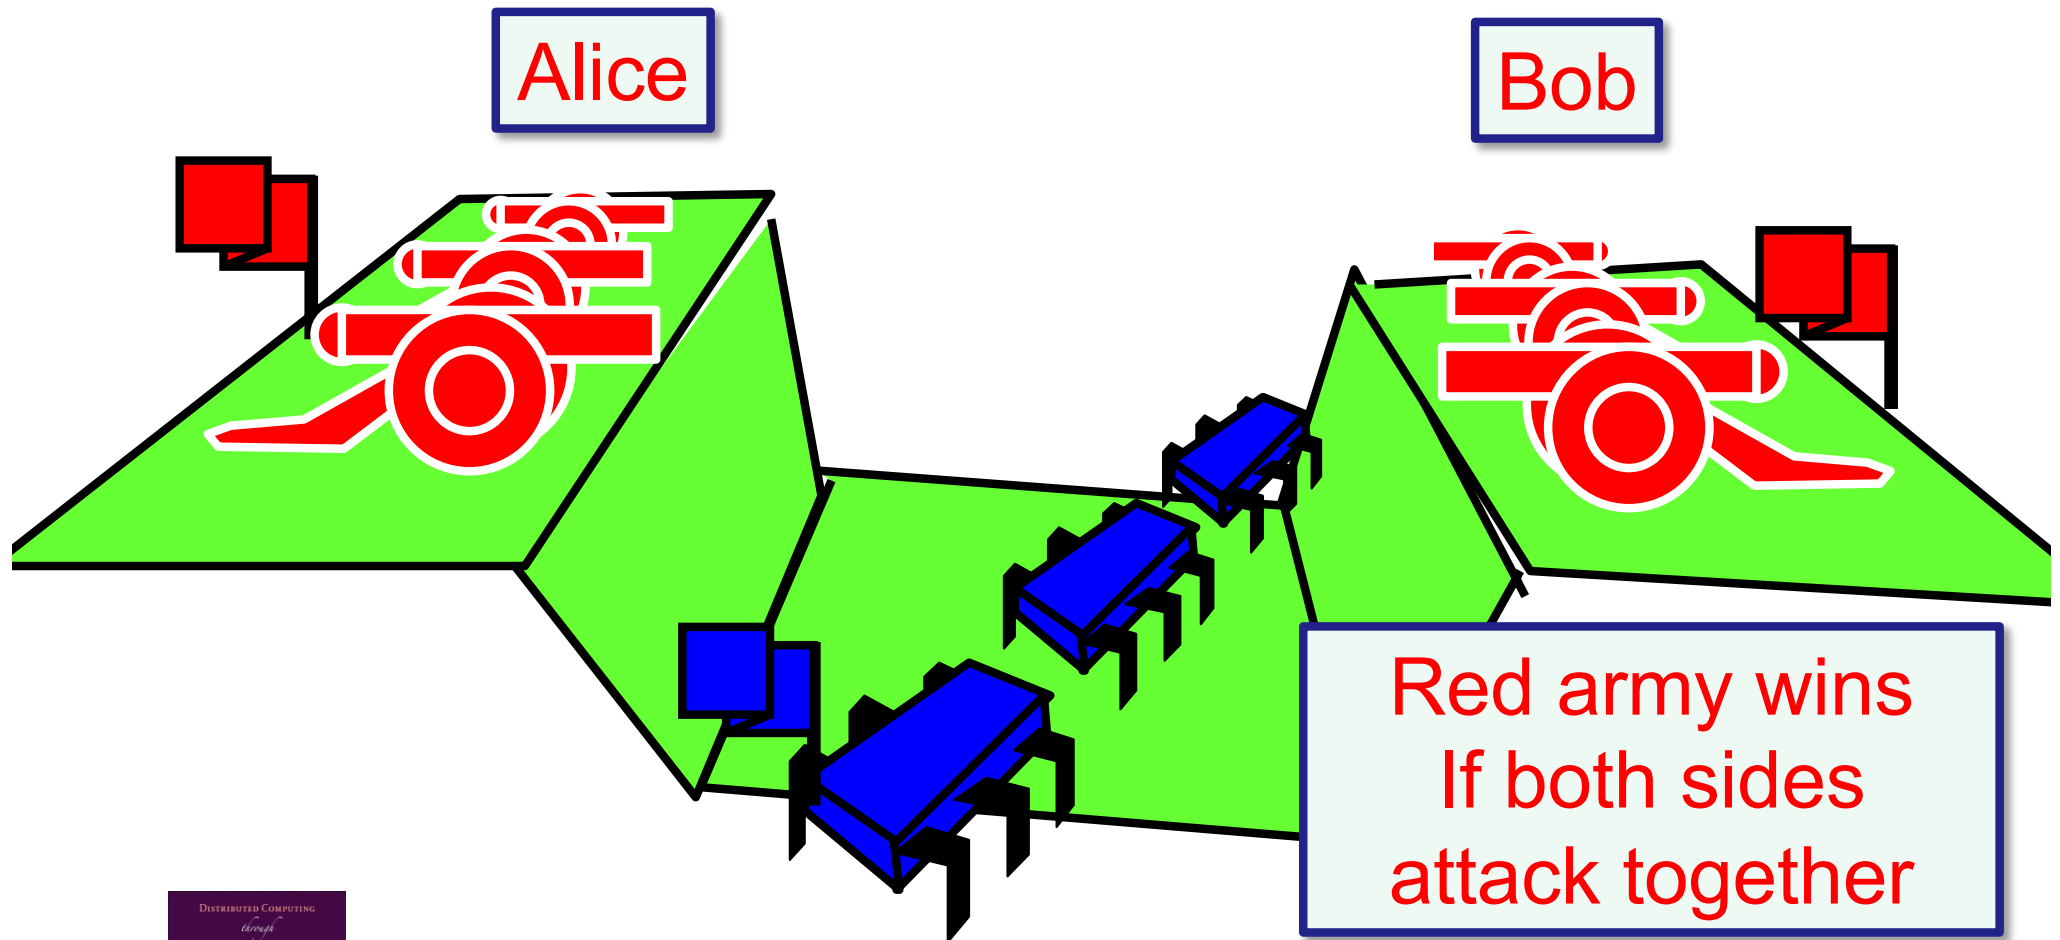
Two Classic Distributed Problems

The Muddy Children

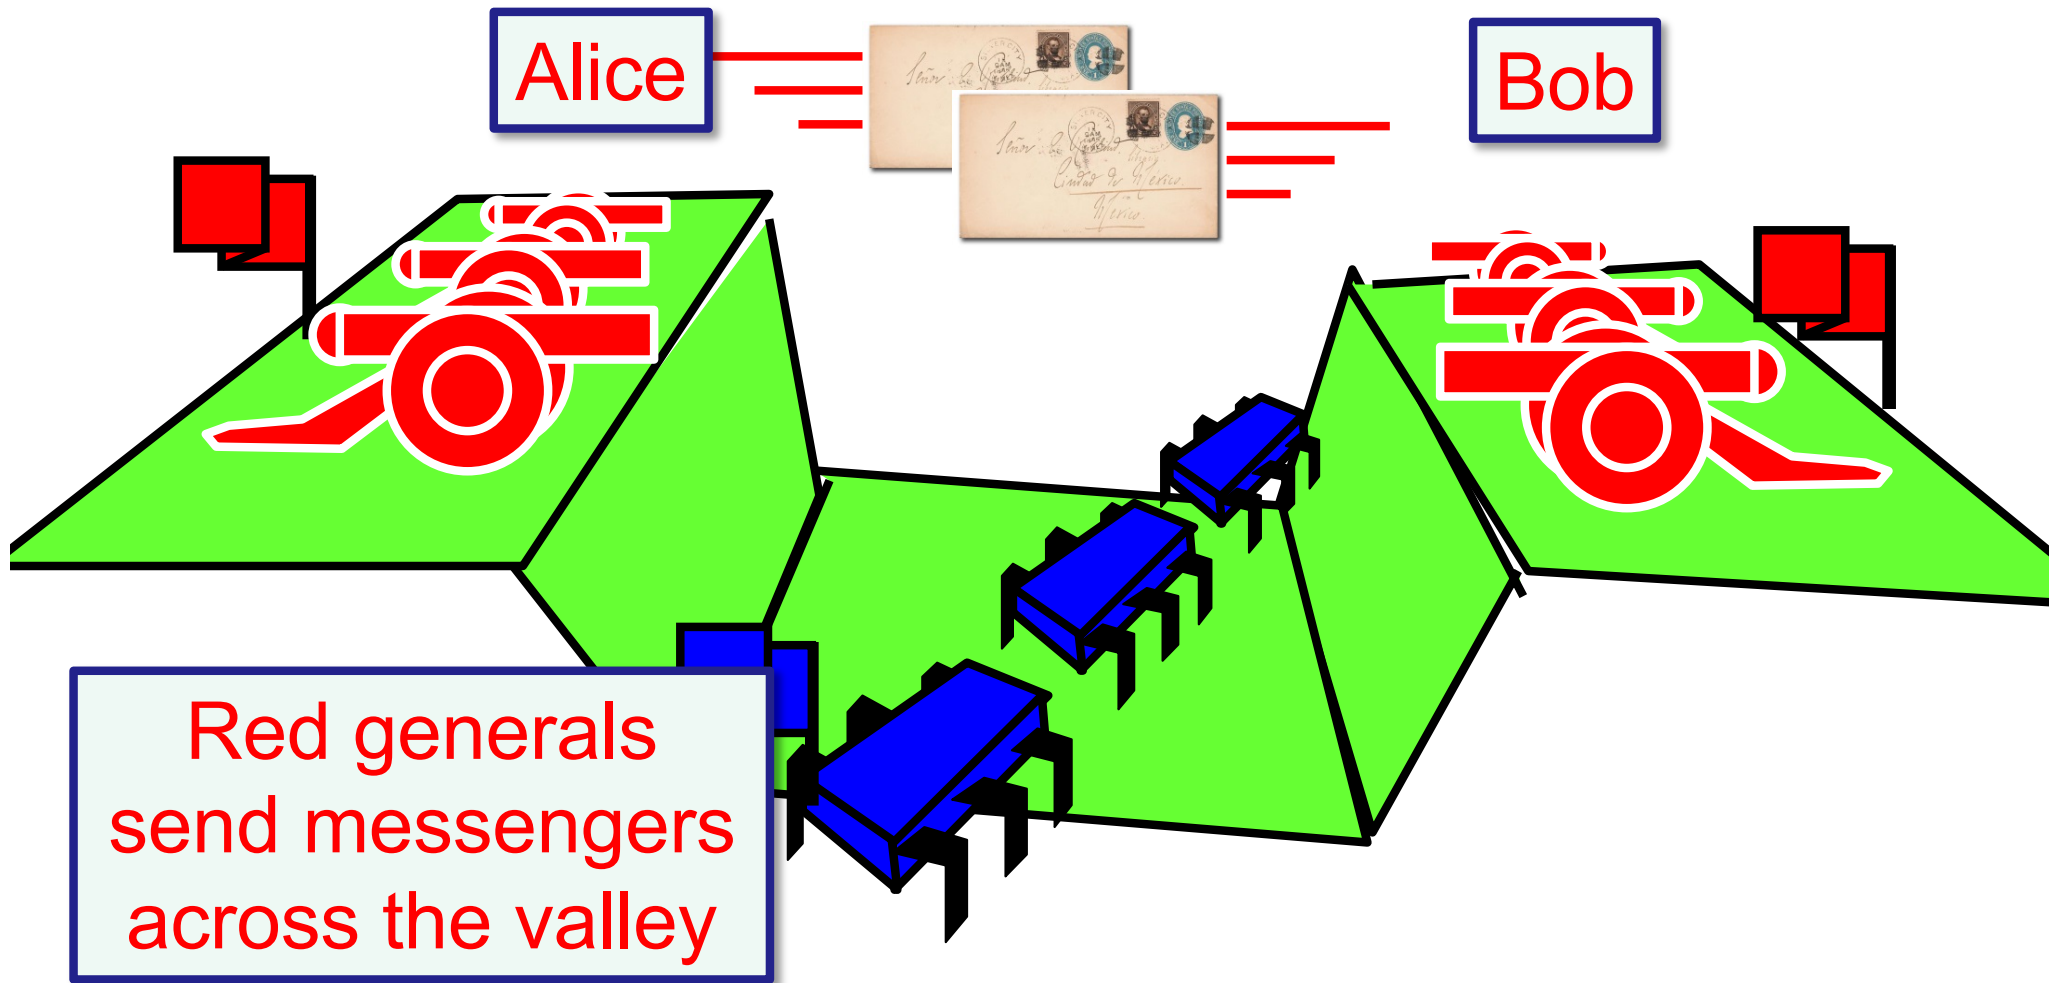
Coordinated Attack



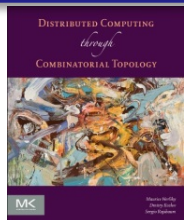
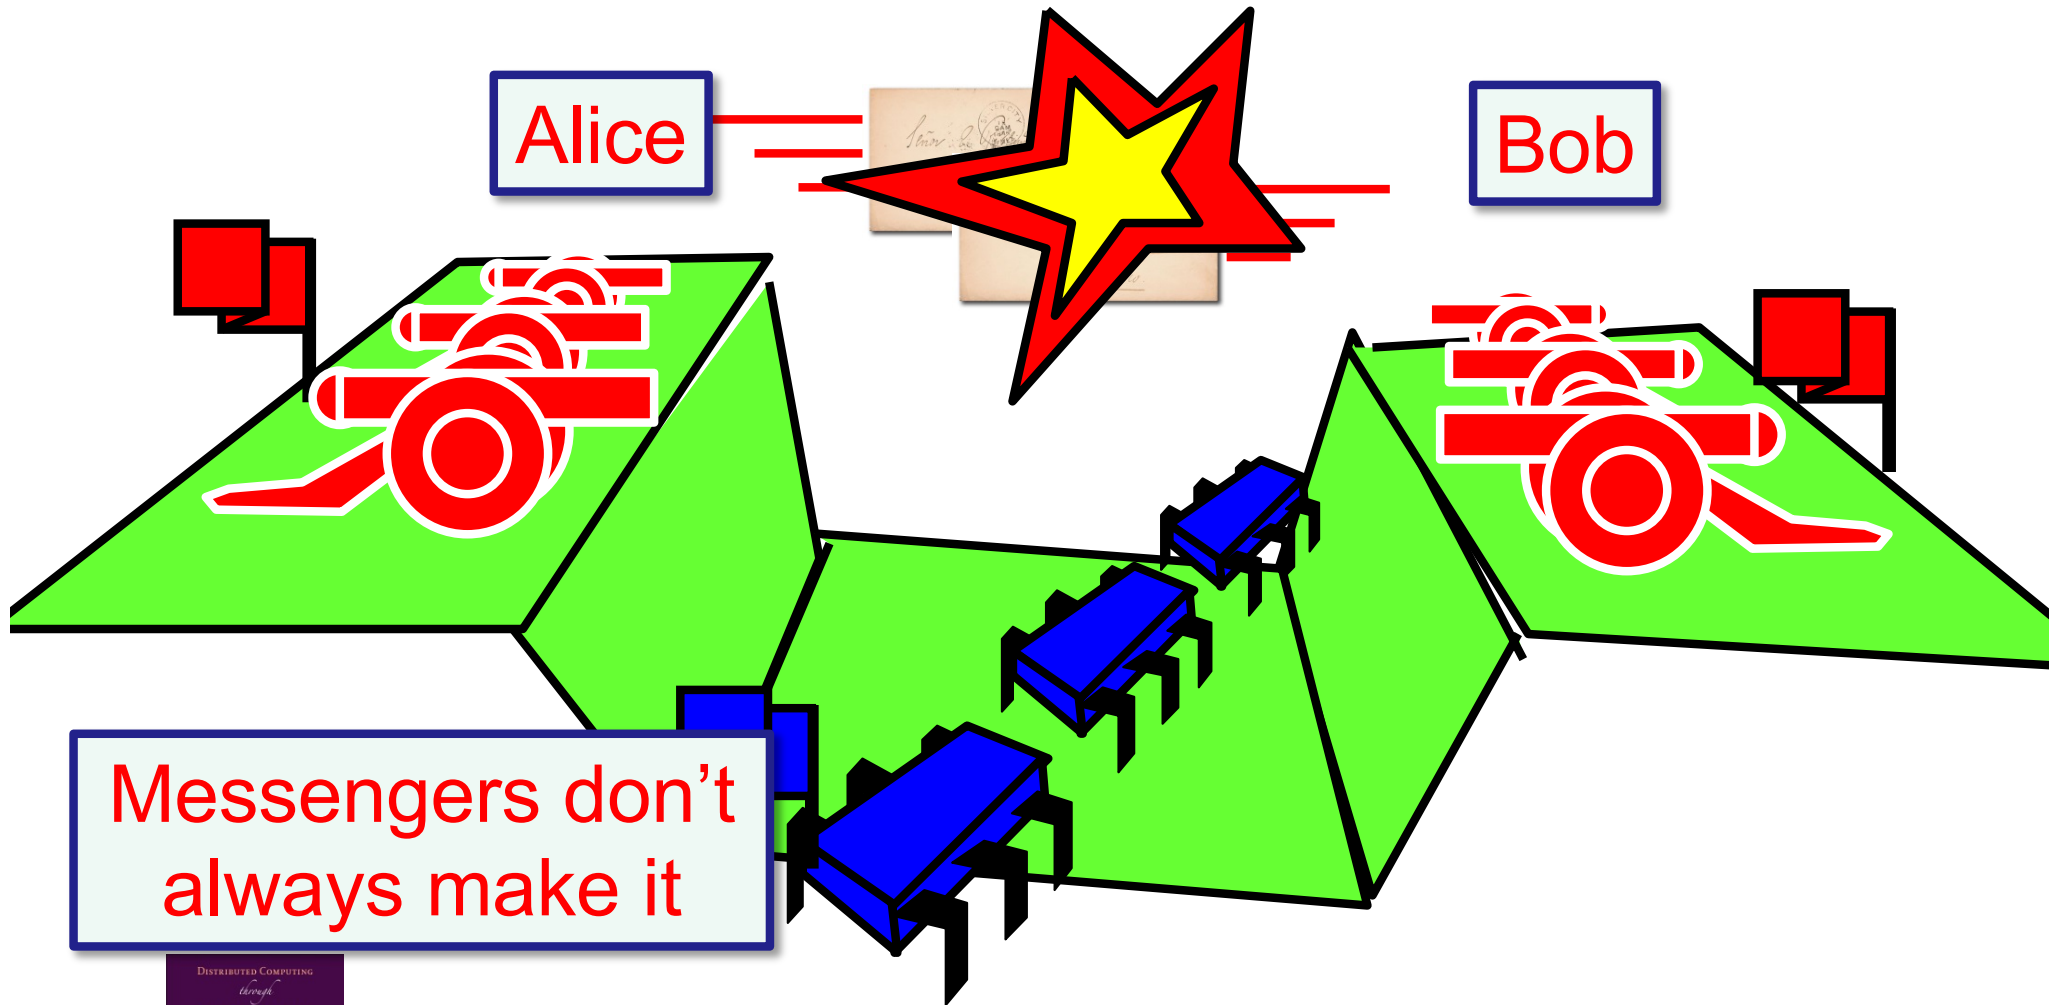
Coordinated Attack



The Two Generals

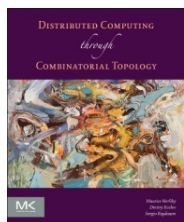


The Two Generals



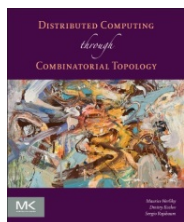
Your Mission

Design a protocol to ensure that
Alice and Bob attack
simultaneously



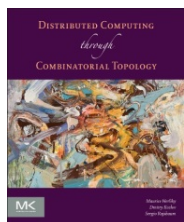
Theorem

There is **no** protocol that ensures that the Red armies attack simultaneously



Operational Proof

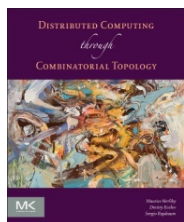
Suppose Bob receives a message at 1:00 saying “attack at Dawn”.



Operational Proof

Suppose Bob receives a message at 1:00 saying “attack at Dawn”.

Are we done?



Operational Proof

Suppose Bob receives a message at 1:00 saying “attack at Dawn”.

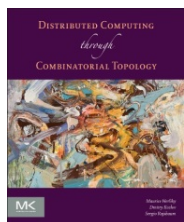
Are we done?

No, because Alice doesn't know if Bob got that message ...

Operational Proof

So Bob sends an
acknowledgment to Alice

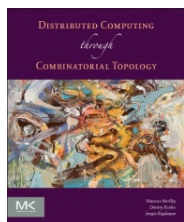
...



Operational Proof

So Bob sends an
acknowledgment to Alice

Are we done?



Operational Proof

So Bob sends an
acknowledgment to Alice

Are we done?

No, because Bob doesn't
know if Alice got that
message ...

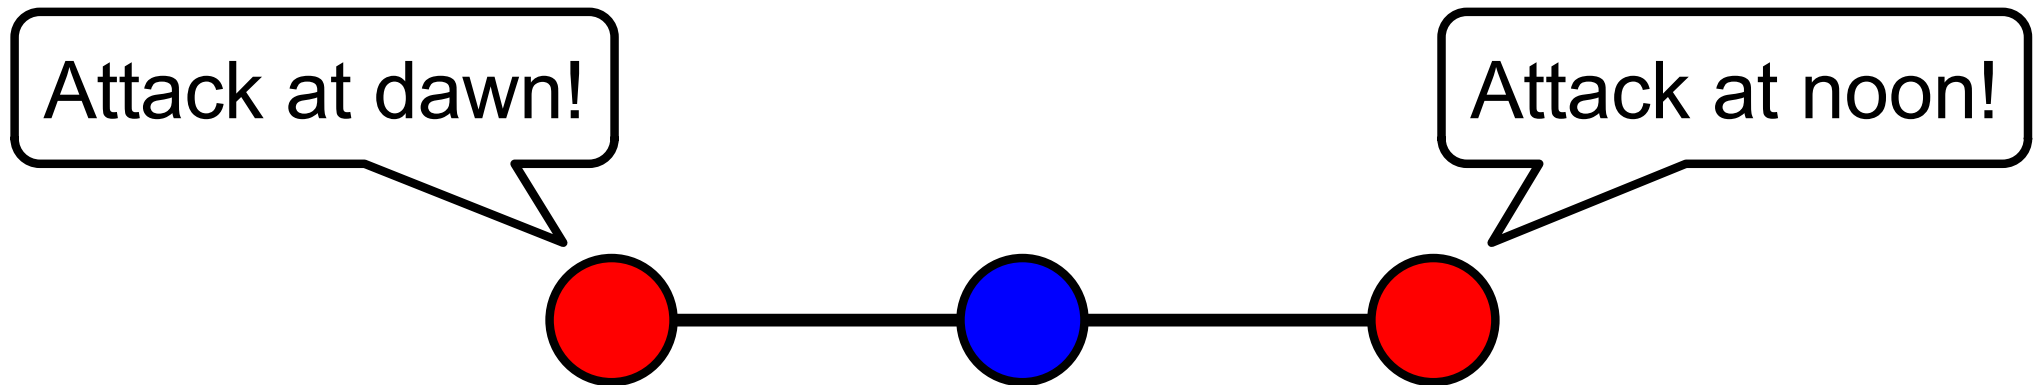
Operational Proof

Bob sends an
acknowledgment to Alice

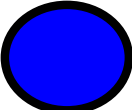
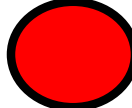
*This goes on forever ... so
no protocol is possible*

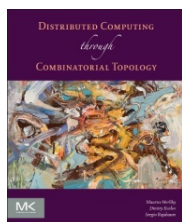
Are we?

No, because Bob does
not know if Alice got that
message ...

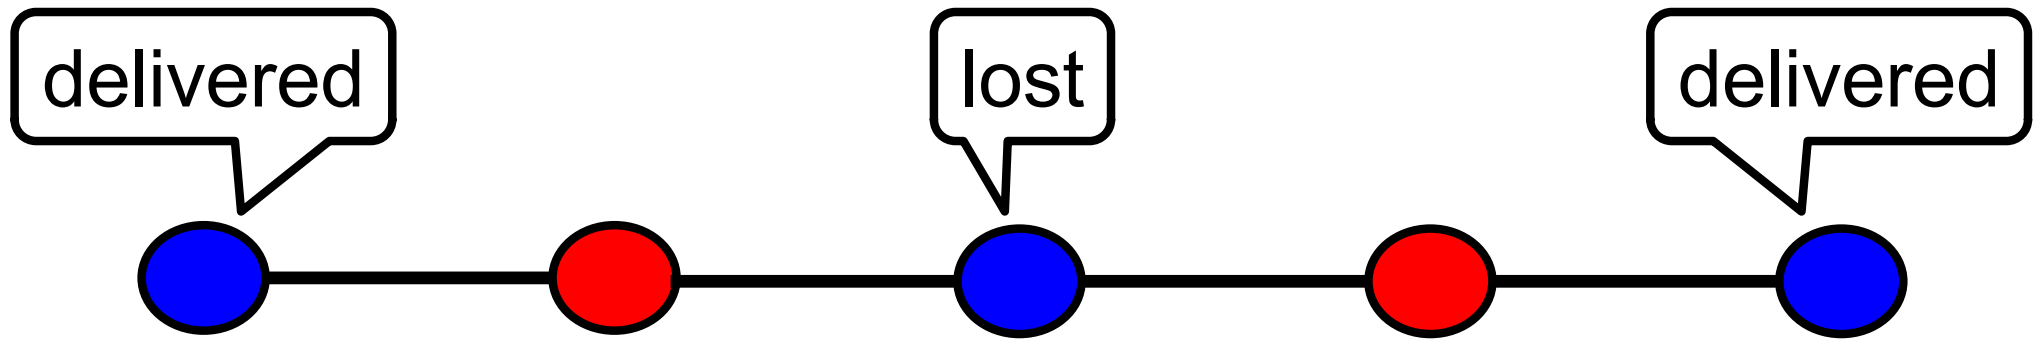


Noon

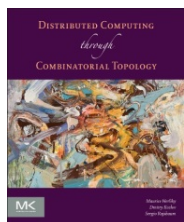
Bob is 
Alice is 



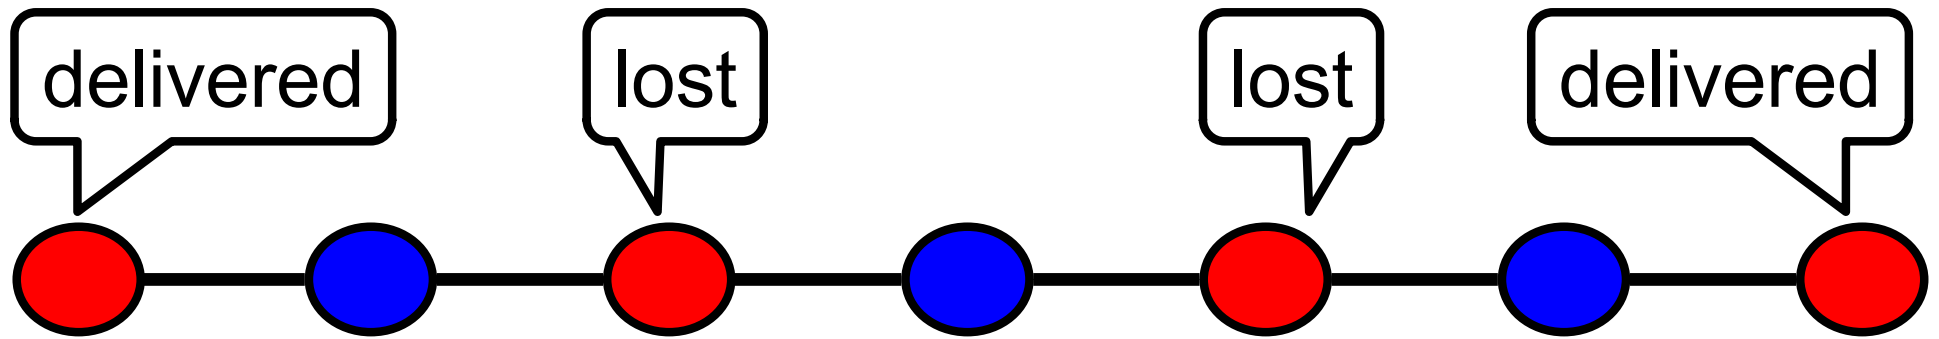
Distributed Computing through
Combinatorial Topology



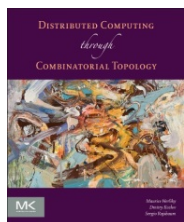
1:00 PM



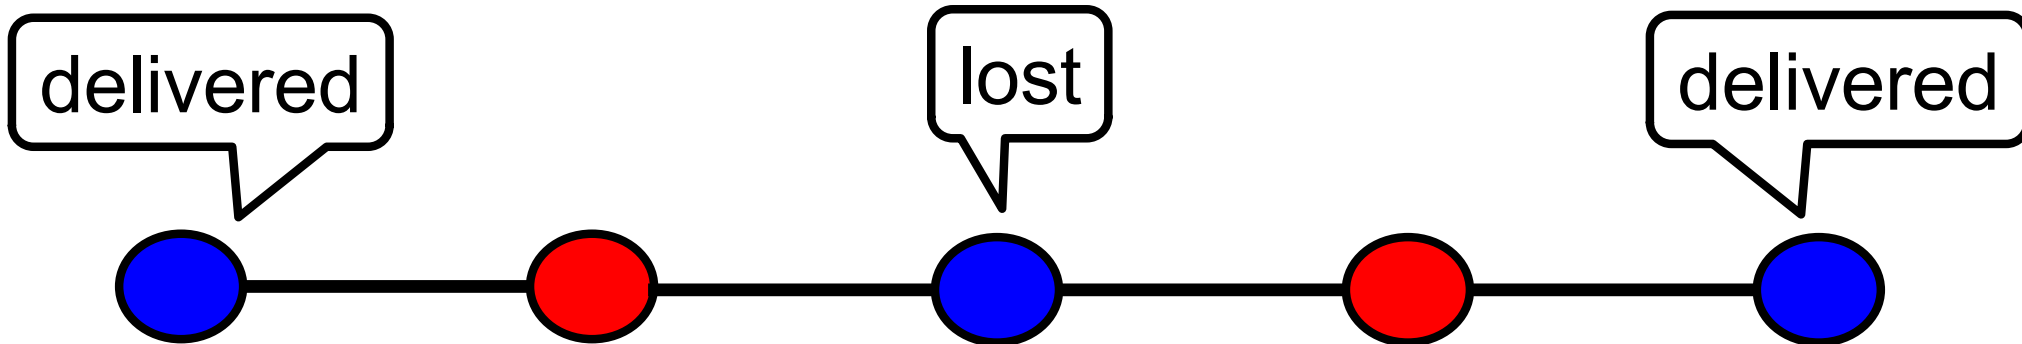
Distributed Computing through
Combinatorial Topology



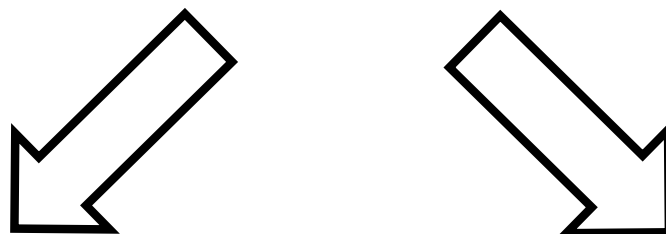
2:00 PM



Distributed Computing through
Combinatorial Topology



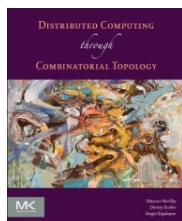
protocol graph

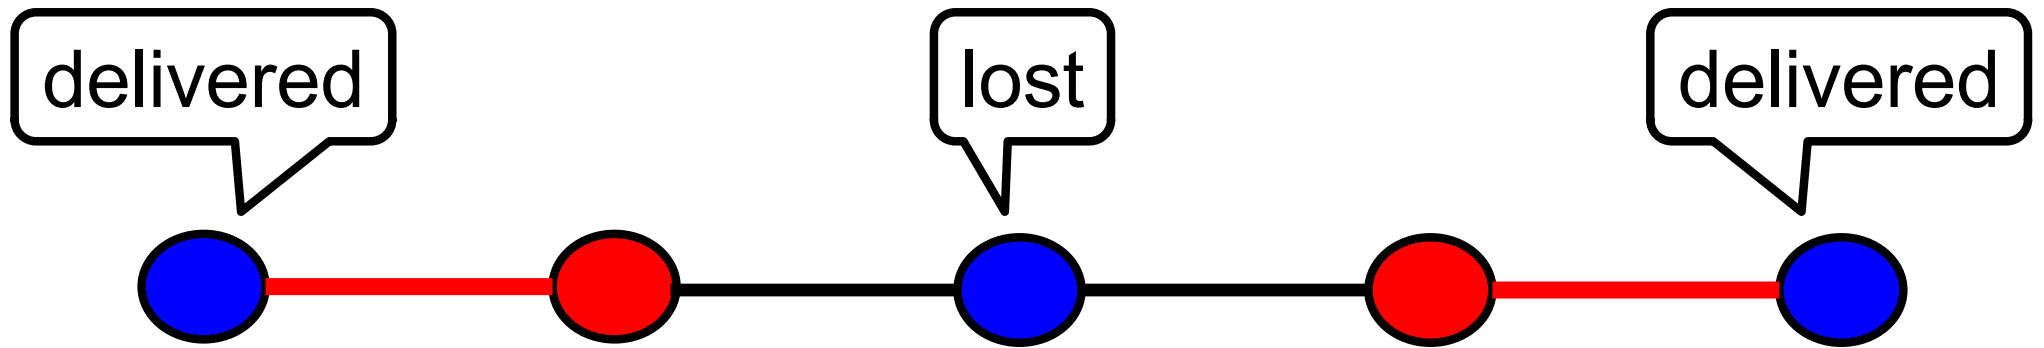


decision
map



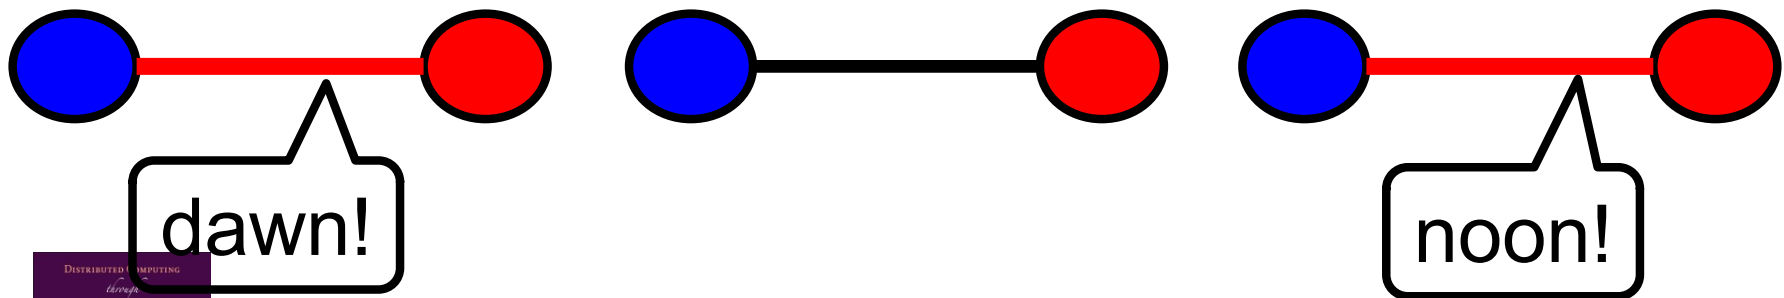
output graph



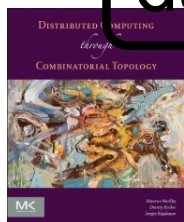


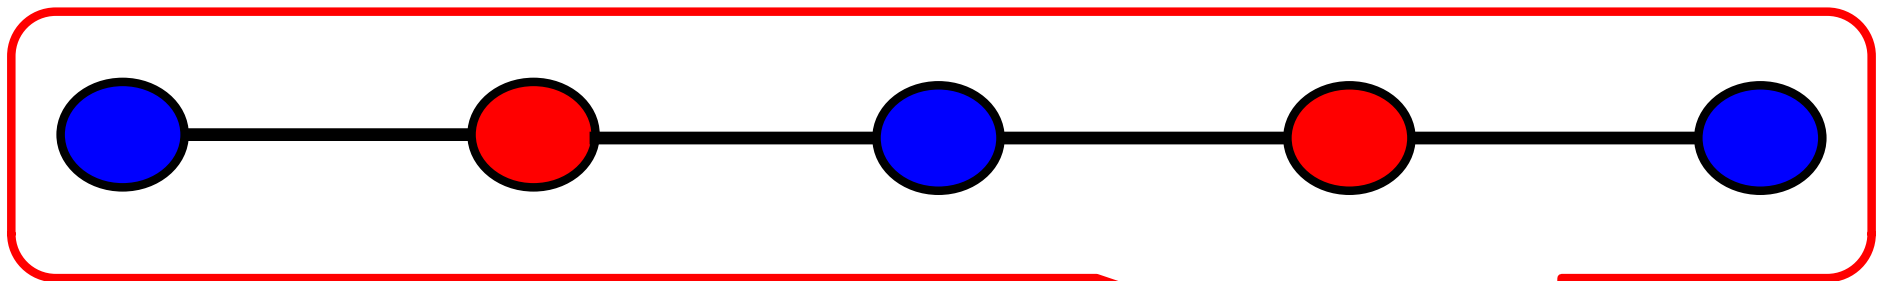
protocol graph

These edges go here

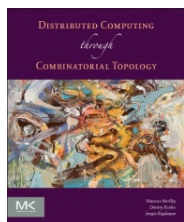


output graph

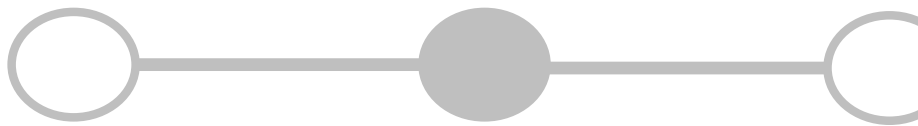




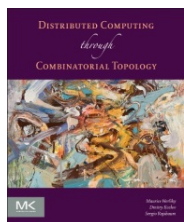
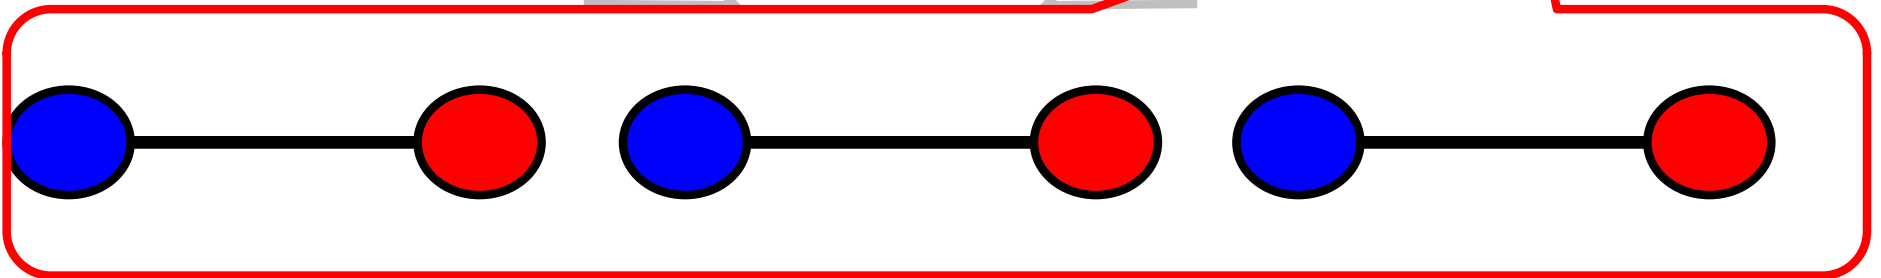
This graph is *connected*



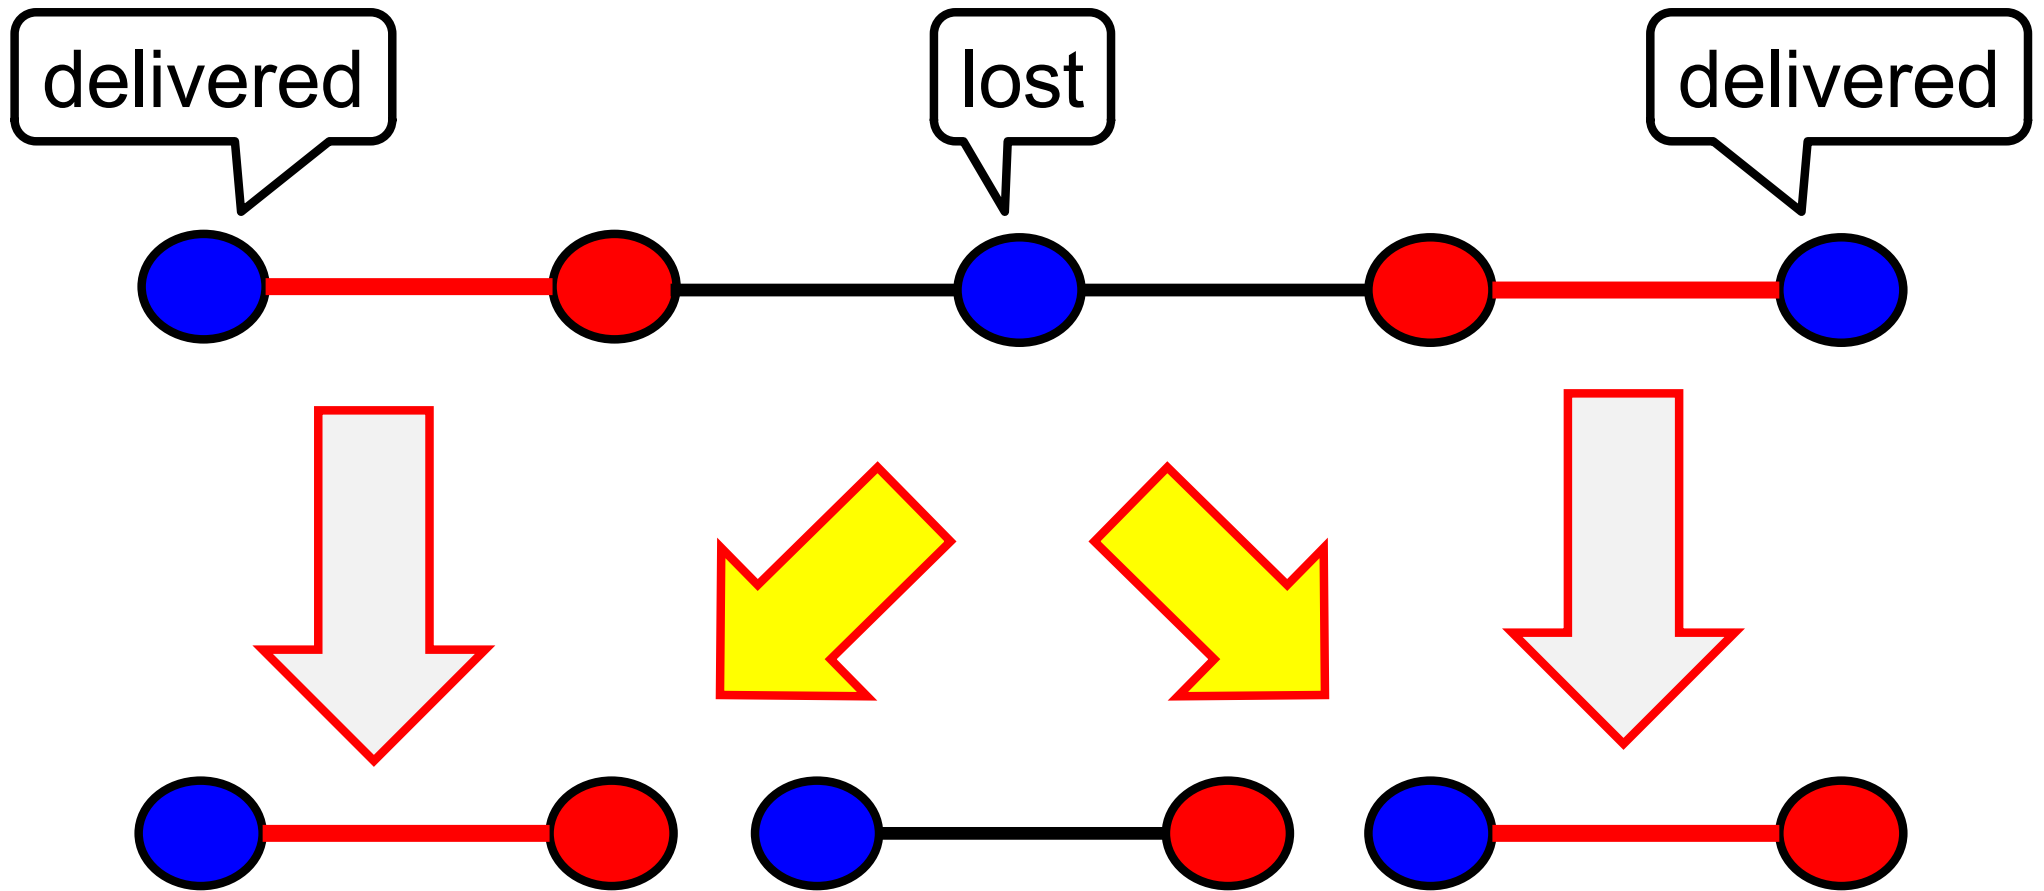
Distributed Computing through
Combinatorial Topology



This graph is
not connected

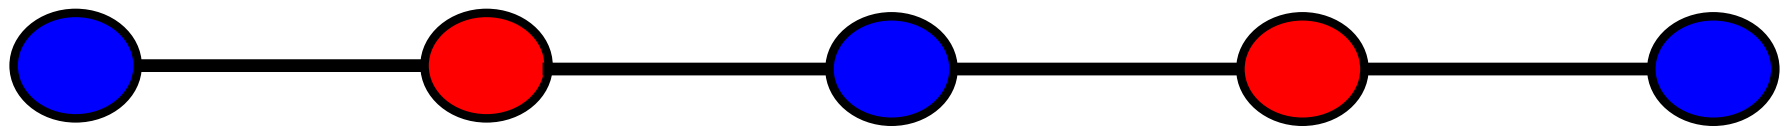


Distributed Computing through
Combinatorial Topology

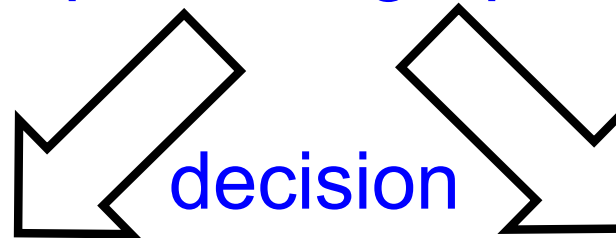


Map not allowed to "tear"
protocol complex





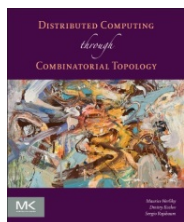
protocol graph



decision
map



output graph

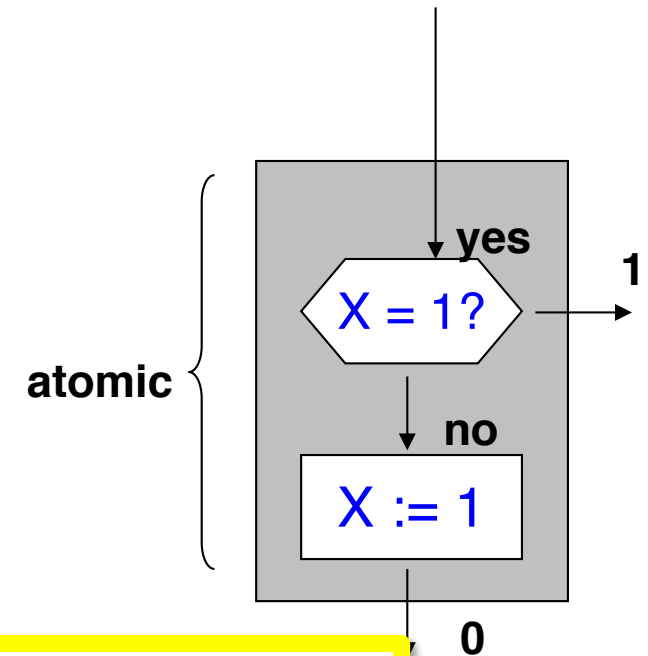


Test and Set

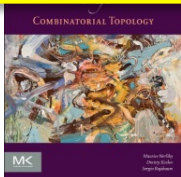
- **TAS(X)** **tests** if $X = 1$, **sets** X to 1 if not, and returns the old value of X
 - Instruction available on almost all processors

TAS(X) :

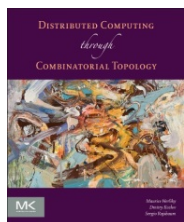
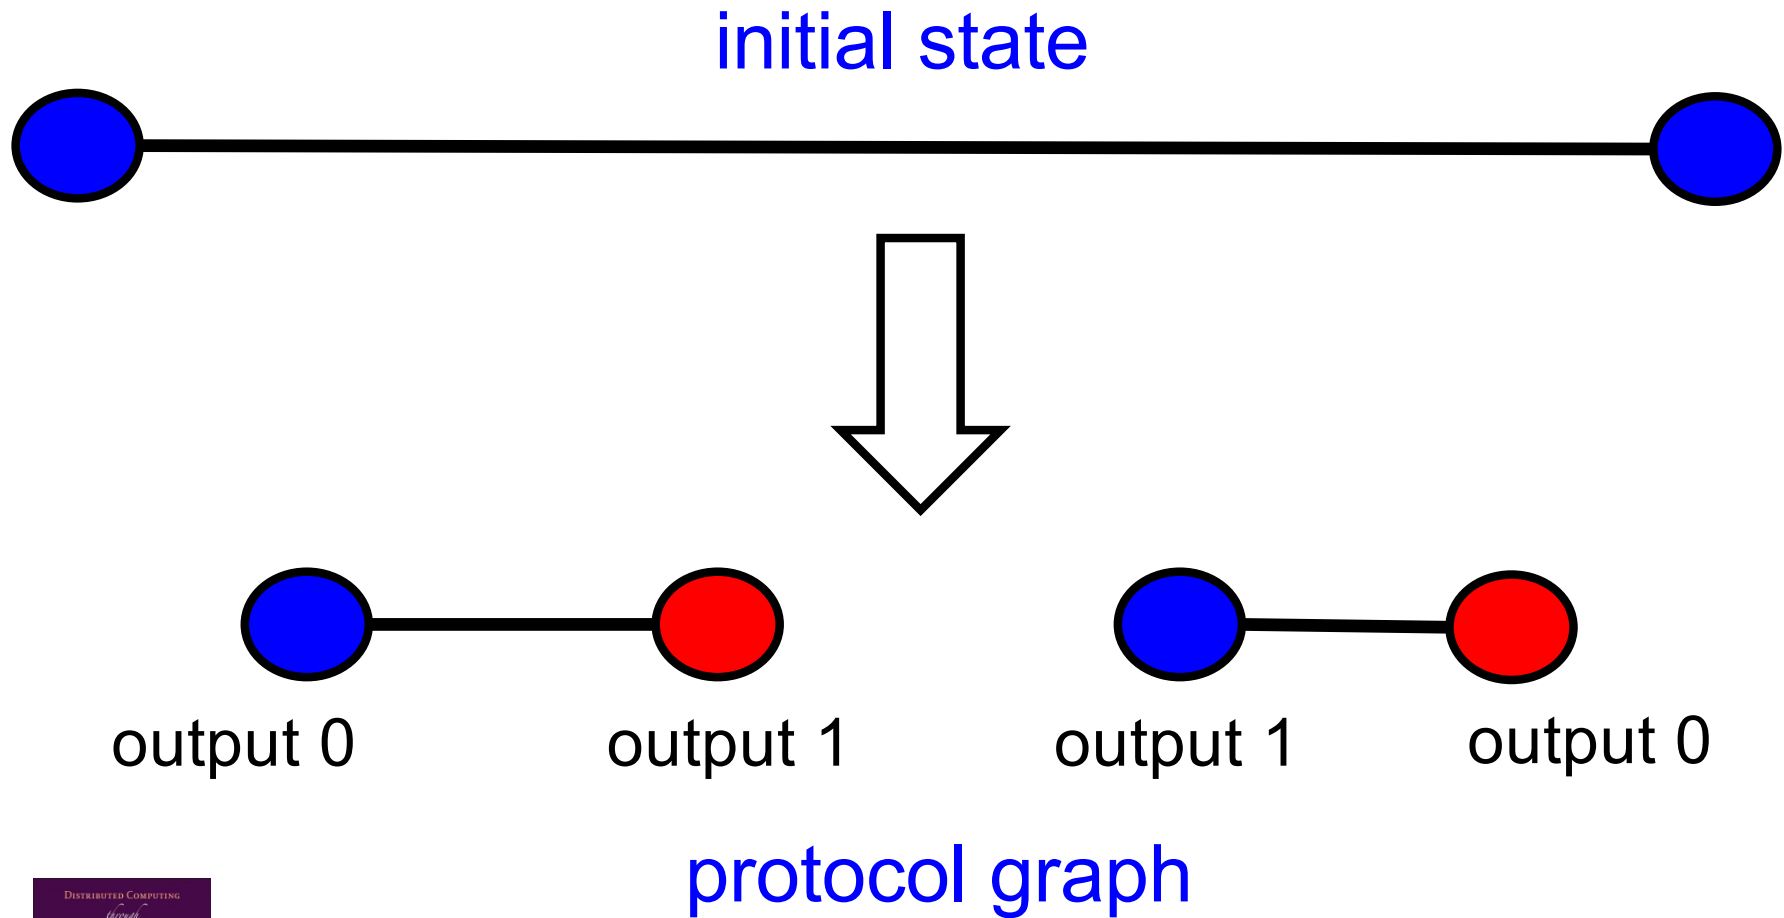
atomic {
 if $X == 1$ return 1;
 $X = 1$;
 return 0;
}



2 processes, P and Q, perform TAS(X)
What is the protocol complex?



2-Process Test and Set



What if there are 3 processes: P, Q, and R?

HW: (Simplified) Peterson's lock: 2 processes

```
bool flag[0]    = false;
bool flag[1]    = false;
int turn;
```

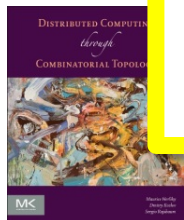
P0:

```
flag[0] = true;
turn = 1;
if(flag[1] and turn==1){
return false // failure
}
return true
// critical section
```

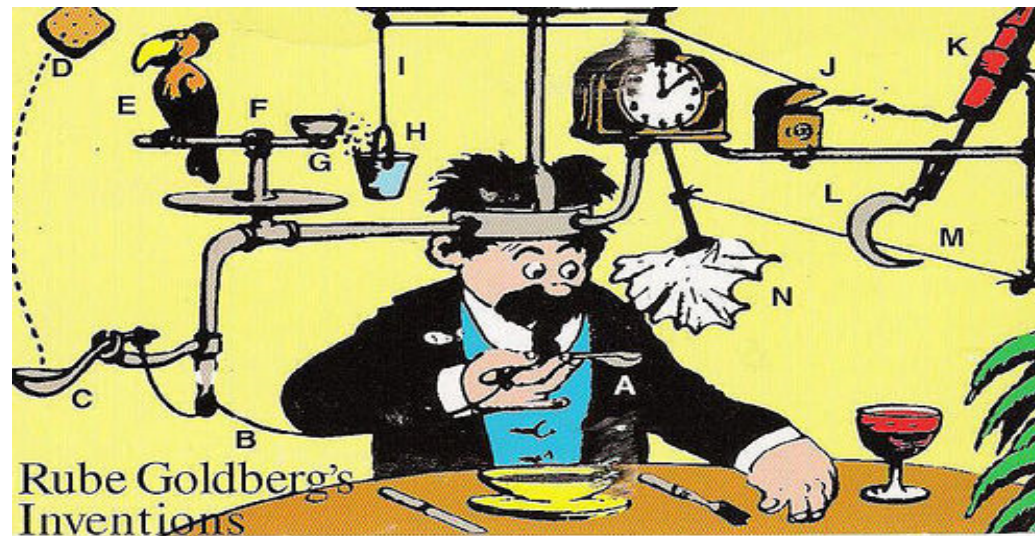
P1:

```
flag[1] = true;
turn = 0;
if(flag[0] and turn==0)
{
return false // failure
}
return true
// critical
section
```

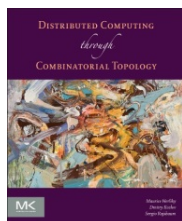
What is the protocol complex?
Can we prove that the two processes cannot
be both in the critical section



Operational Reasoning

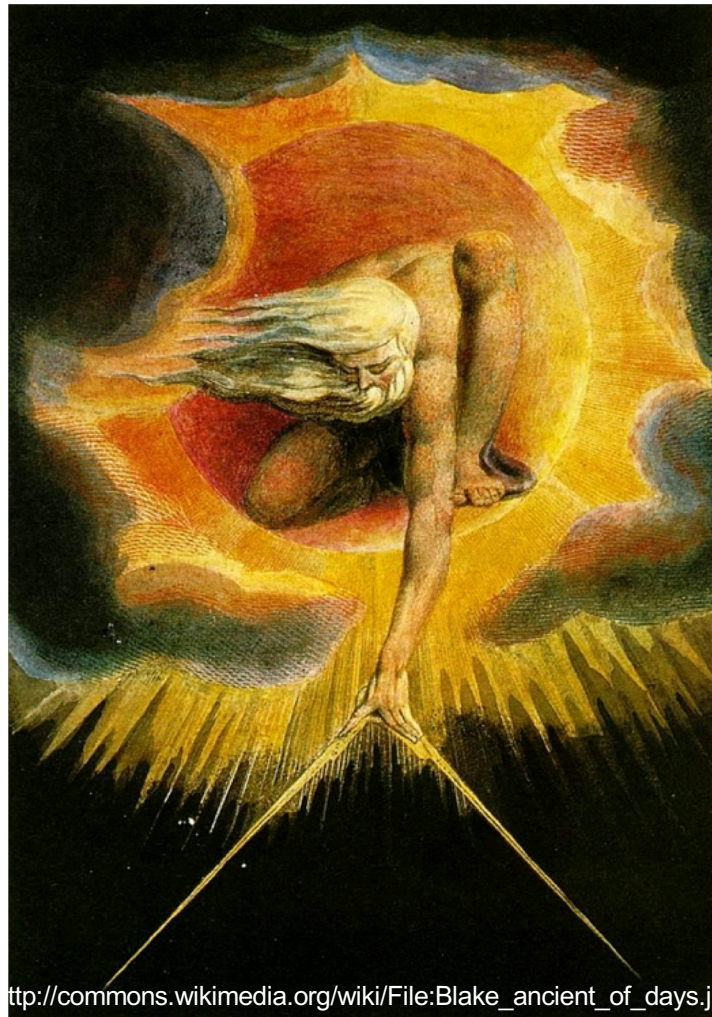


http://commons.wikimedia.org/wiki/File:Professor_Lucifer_Butts.gif

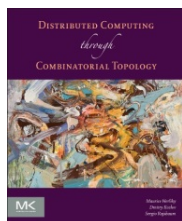


Distributed Computing through
Combinatorial Topology

Combinatorial Reasoning

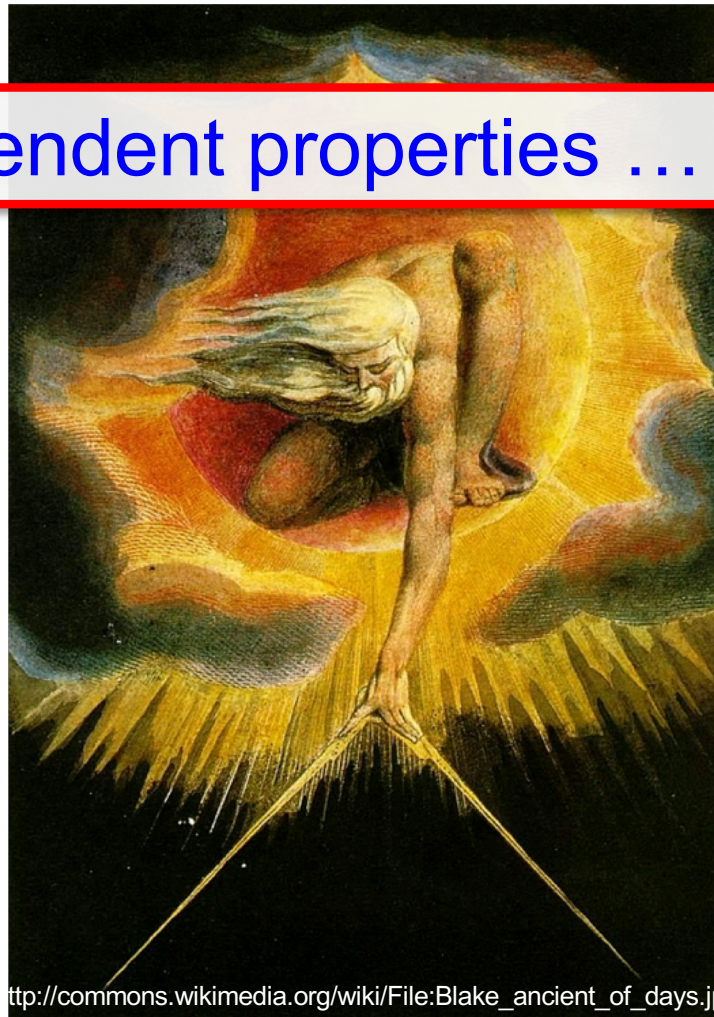


http://commons.wikimedia.org/wiki/File:Blake_ancient_of_days.jp

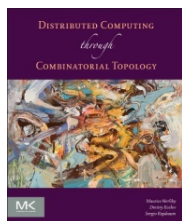


Combinatorial Reasoning

Model-independent properties ...



http://commons.wikimedia.org/wiki/File:Blake_ancient_of_days.jp

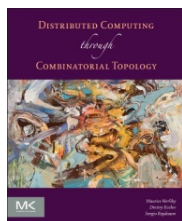


Combinatorial Reasoning

Model-independent properties ...



... restricted model-dependent reasoning

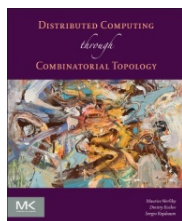


http://commons.wikimedia.org/wiki/File:Blake_ancient_of_days.jp



This work is licensed under a [Creative Commons Attribution-ShareAlike 2.5 License](https://creativecommons.org/licenses/by-sa/2.5/).

- **You are free:**
 - **to Share** — to copy, distribute and transmit the work
 - **to Remix** — to adapt the work
- **Under the following conditions:**
 - **Attribution.** You must attribute the work to “**Distributed Computing Through Combinatorial Topology**” (but not in any way that suggests that the authors endorse you or your use of the work).
 - **Share Alike.** If you alter, transform, or build upon this work, you may distribute the resulting work only under the same, similar or a compatible license.
- For any reuse or distribution, you must make clear to others the license terms of this work. The best way to do this is with a link to
 - <http://creativecommons.org/licenses/by-sa/3.0/>.
- Any of the above conditions can be waived if you get permission from the copyright holder.
- Nothing in this license impairs or restricts the author's moral rights.



Distributed Computing through
Combinatorial Topology