Weak Models For Distributed Computing

Gadi Taubenfeld

IDC, Israel

Weak Models For Distributed Computing

Gadi Taubenfeld

IDC, Israel

Gadi: I am not an implementor of tools, programming languages, or platforms!

Annie: … pls mention where computers can help you except from text editor and a slides editor …

New 2019

Part I

Genome-Wide Epigenetic Modifications as a Shared Memory Consensus Problem

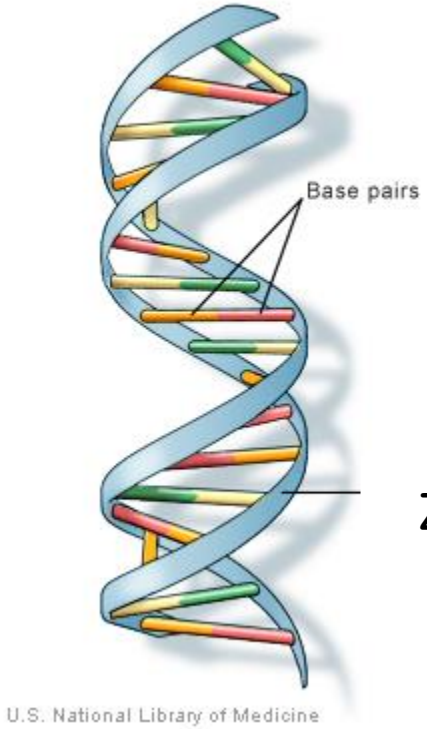

Ziv Bar-Joseph
CMU

Sabrina Rashid
CMU

Gadi Taubenfeld
IDC

The human genome

The entire DNA of a single human cell

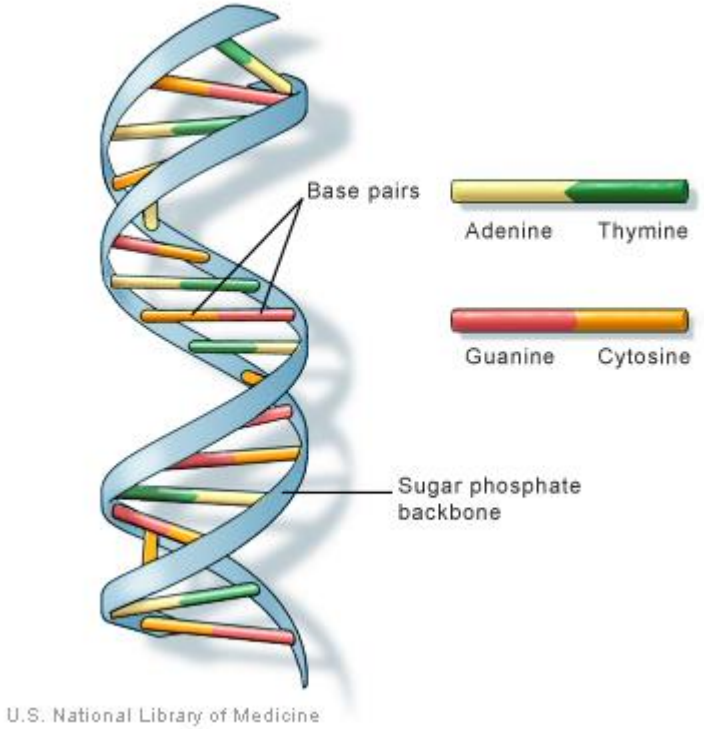

- Two meters long
- 3 billion base pairs
- About 25,000 genes

(Only about 1 percent of DNA is made up of protein-coding genes)

Chromatin

Package DNA into a small volume to fit into the nucleus of a cell

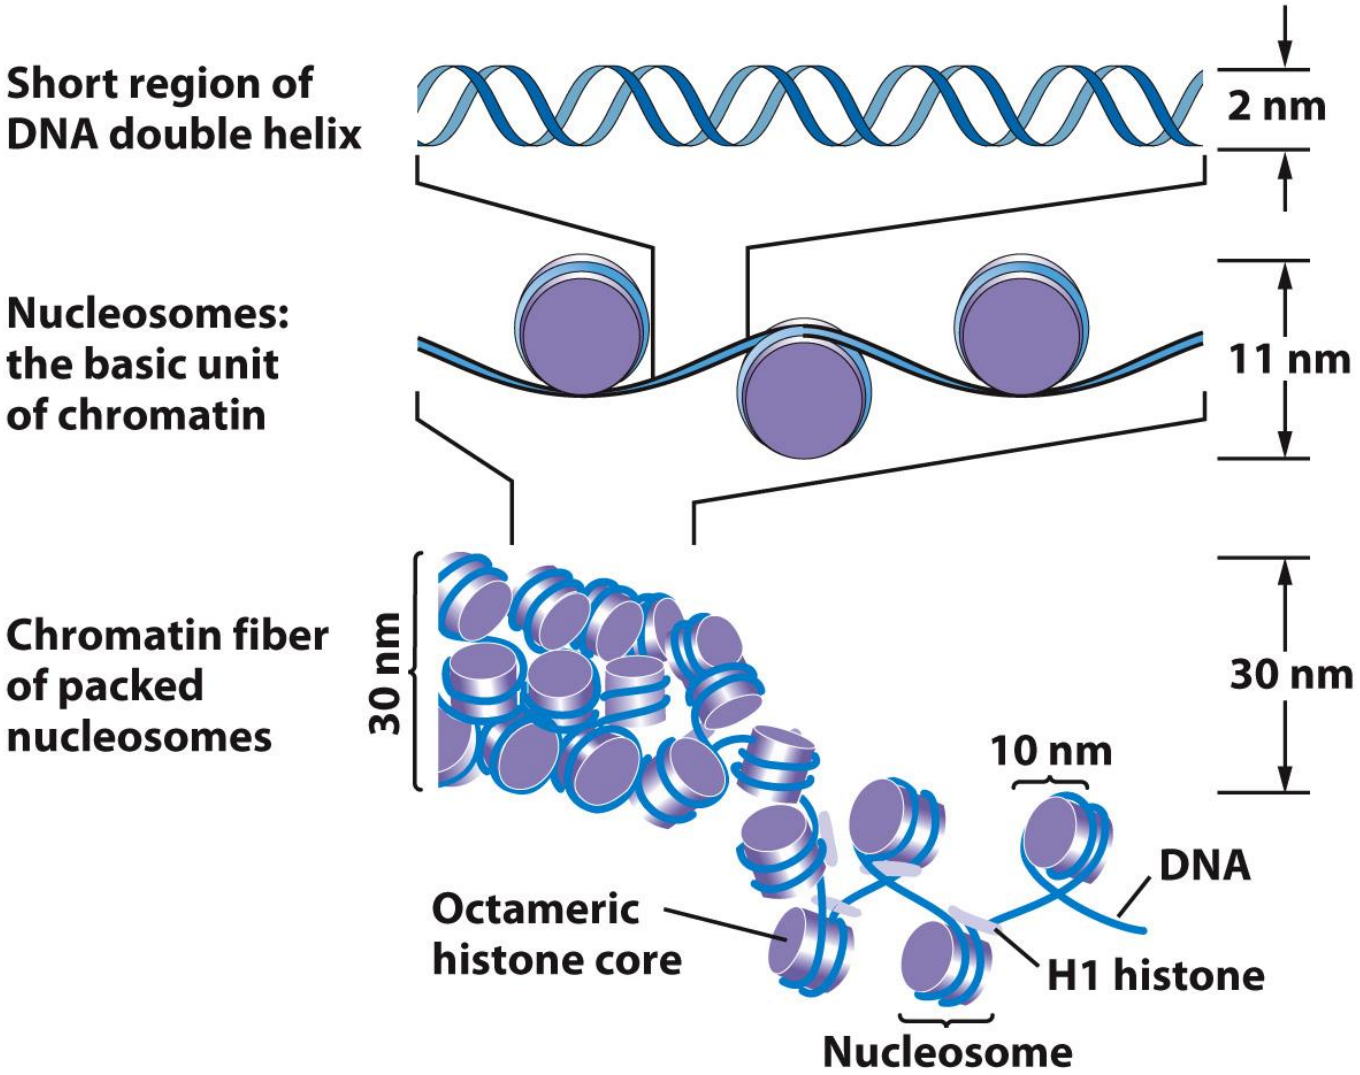

Figure 12-11a
Introduction to Genetic Analysis, Eleventh Edition

Cells types & DNA

Q: How can an organism have different cell types yet one genome?

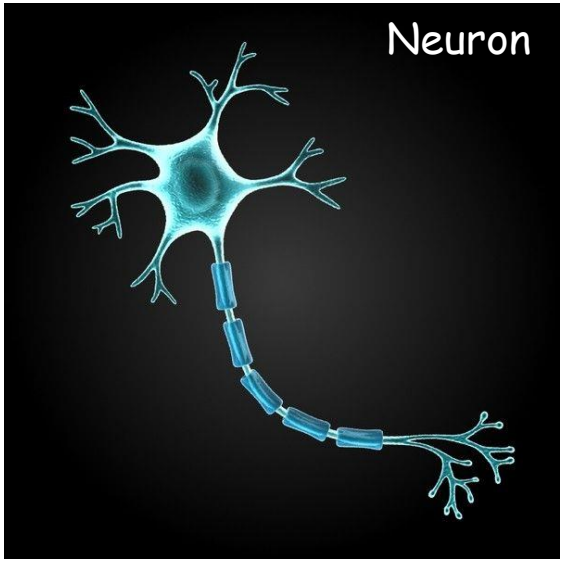

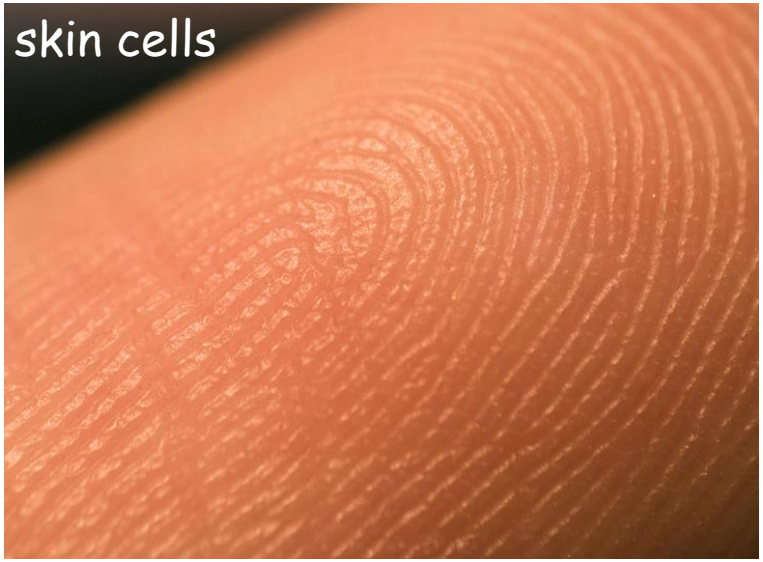

A: Each cell expresses, or turns on, only a fraction of its genes.
The rest of the genes are repressed, or turned off.

Regulation of gene expression

Turning genes on and off

Condensed chromatin

Open chromatin

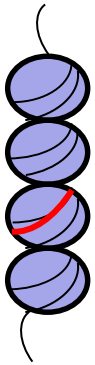

Activate

Deactivate

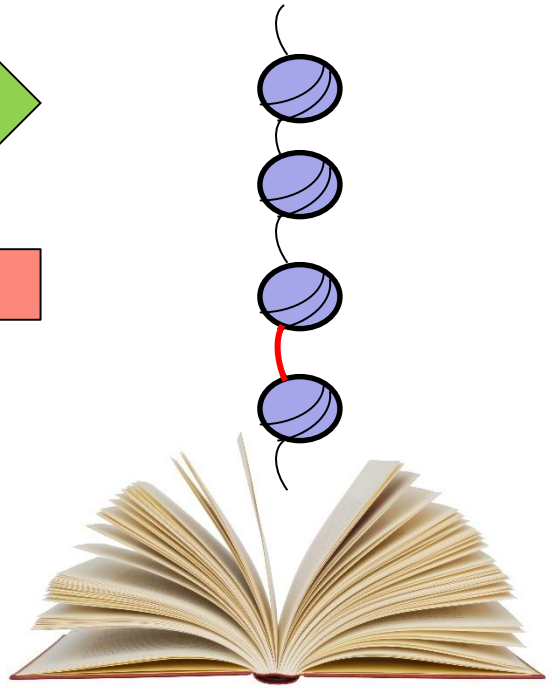

Off

On

Environmental influences, such as a person's diet, stress and exposure to pollutants, impact gene expression.

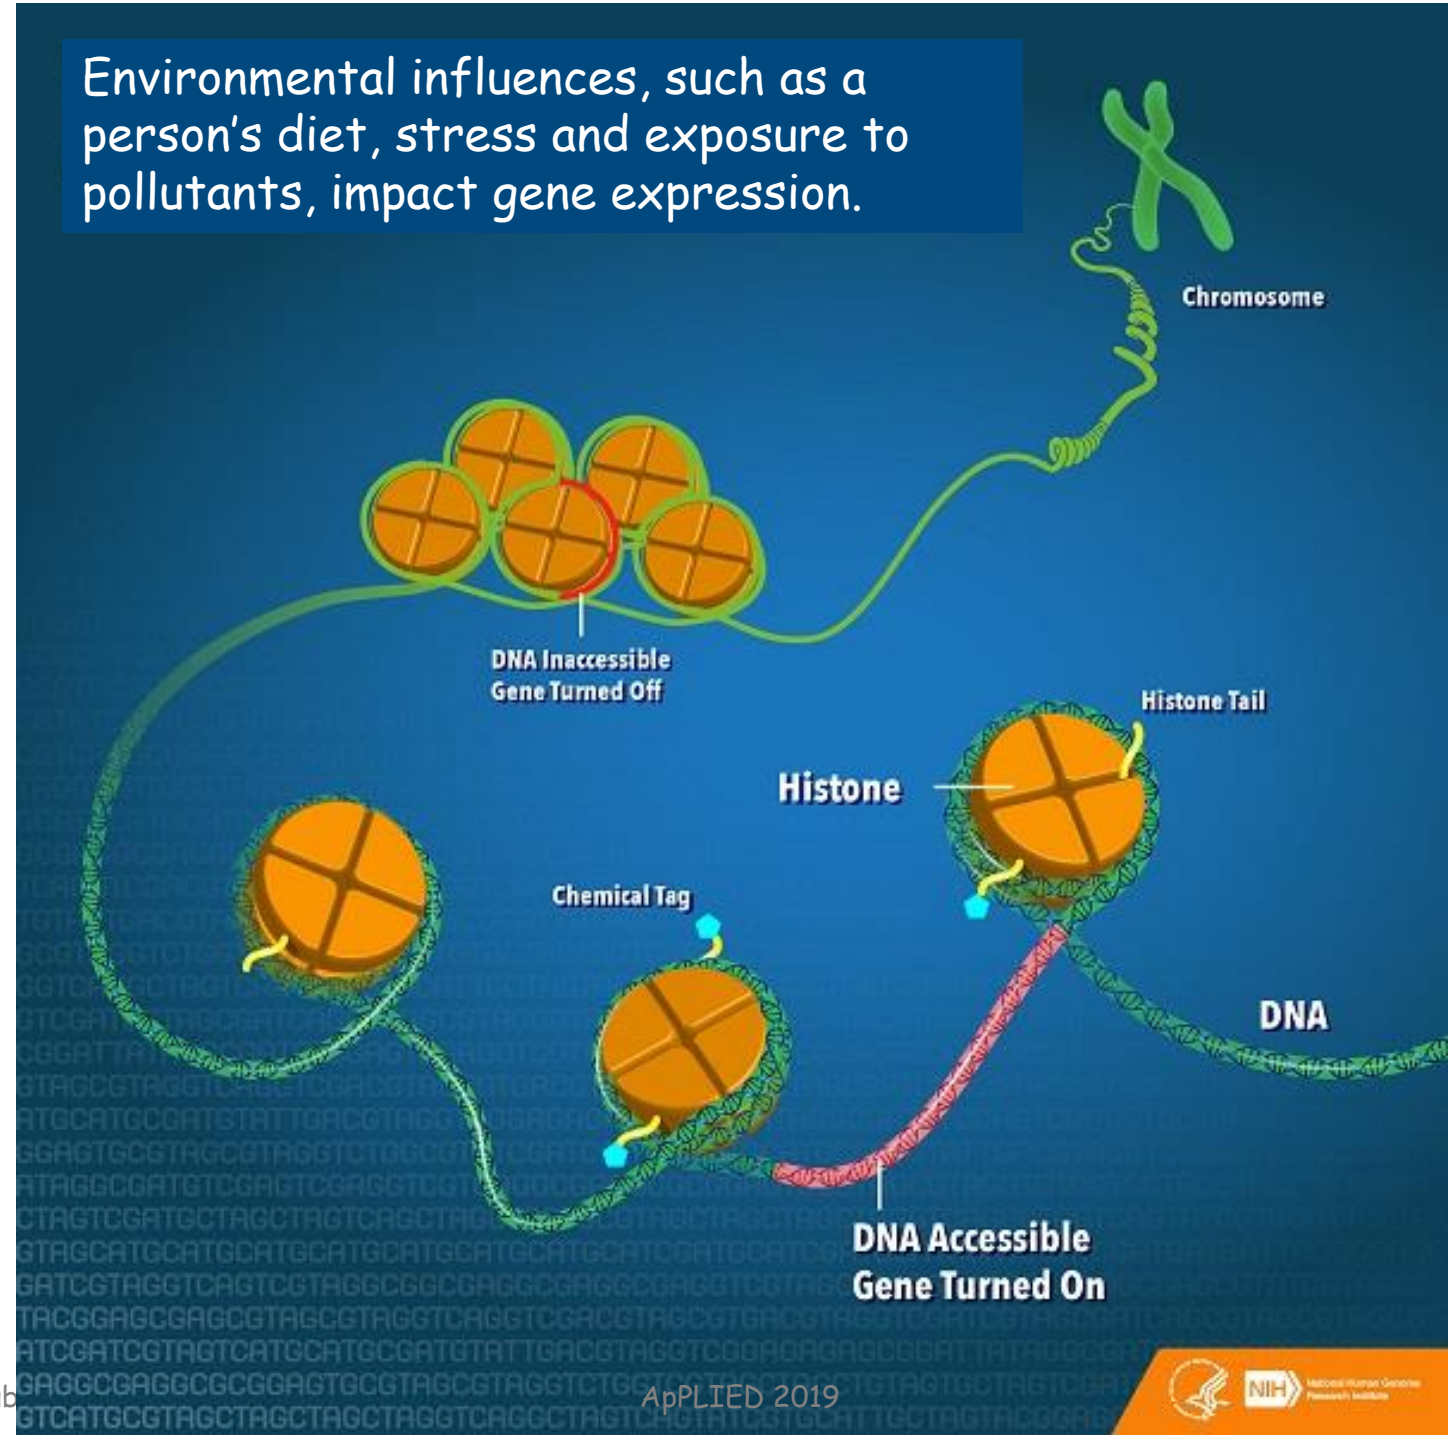

Epigenetics

Modifications that do not change the DNA and affect gene activity

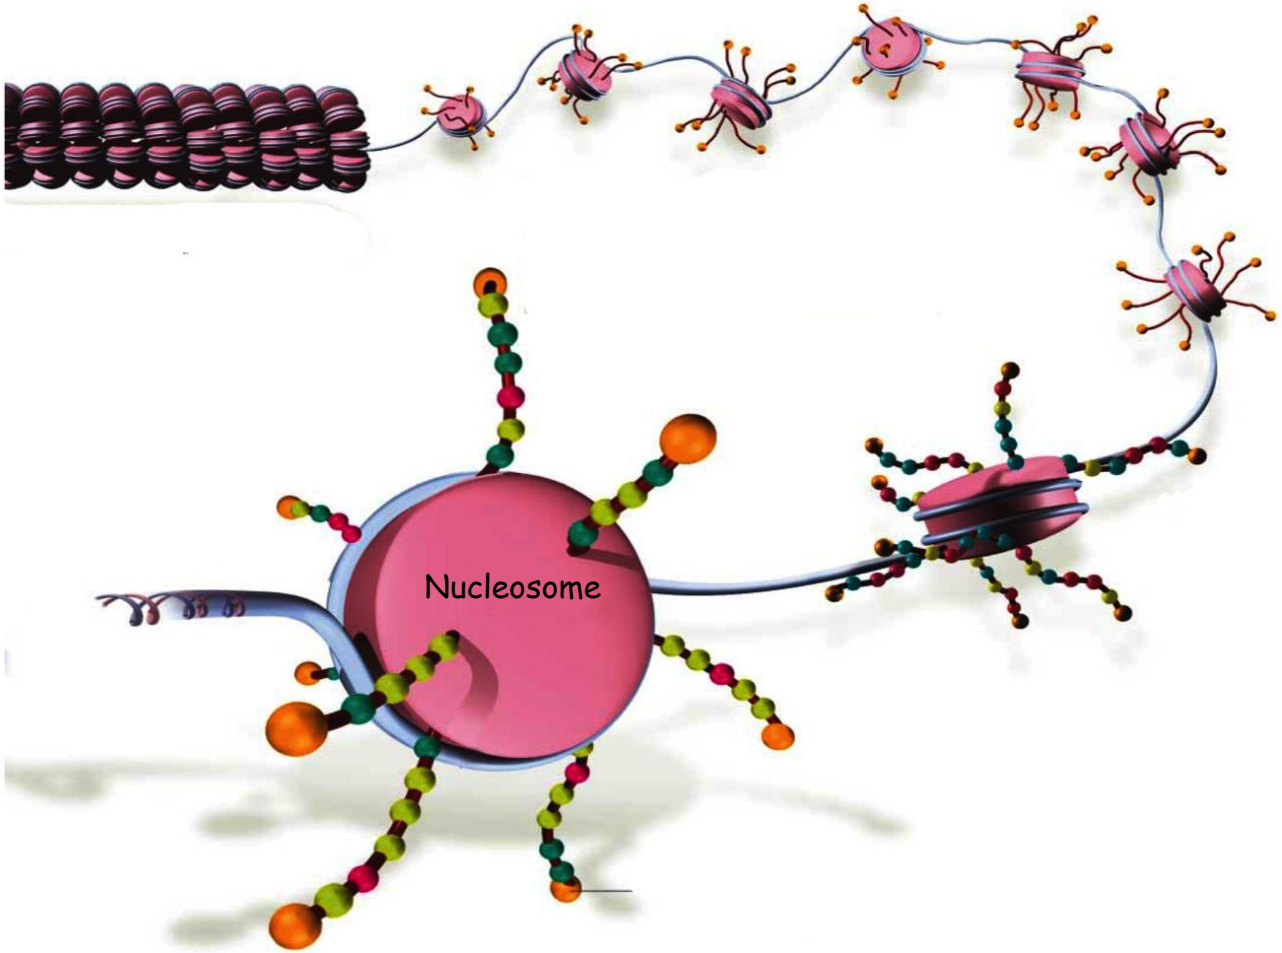

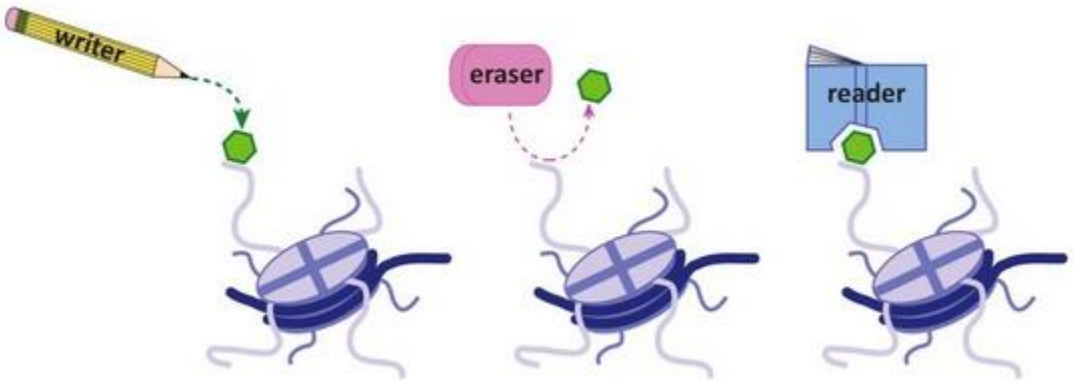
writer
eraser
reader

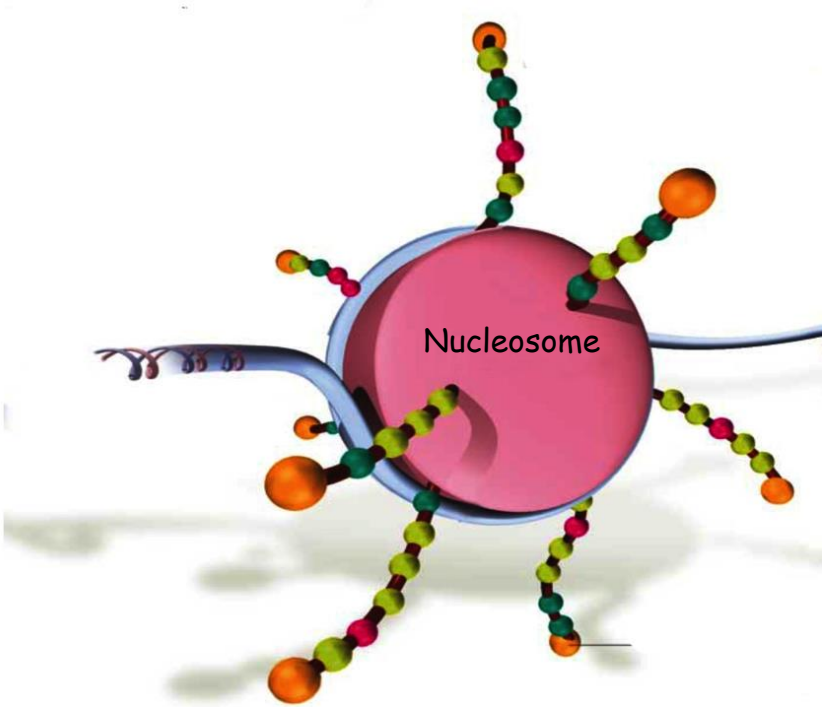
Nucleosome

- 1-writer
- 0-eraser
- 0-writer
- 1-eraser

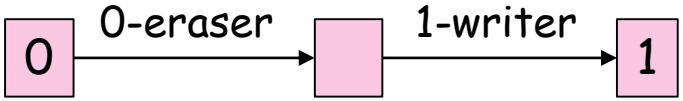

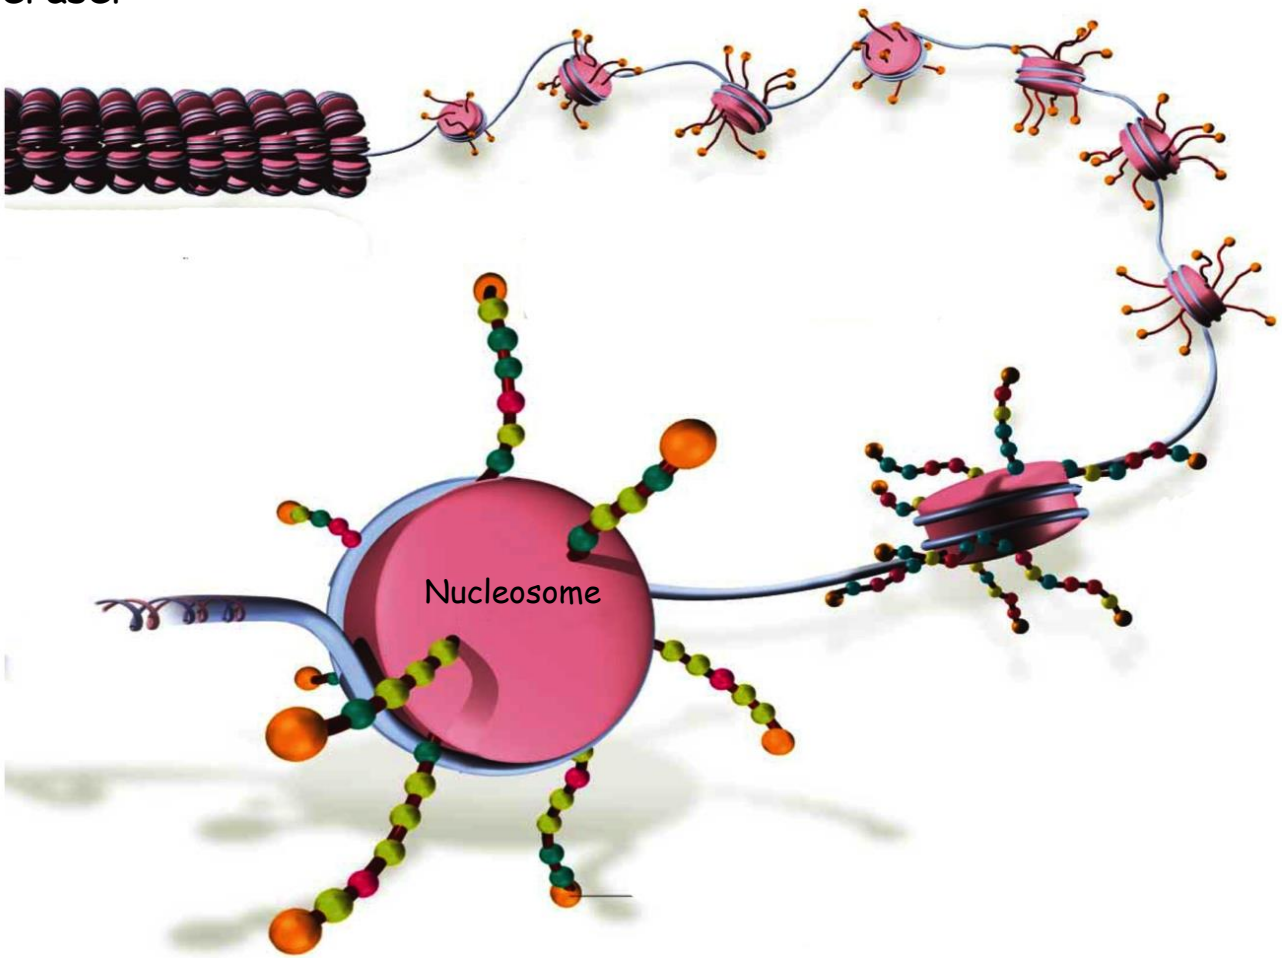

- 1-writer
- 0-eraser
- 0-writer
- 1-eraser

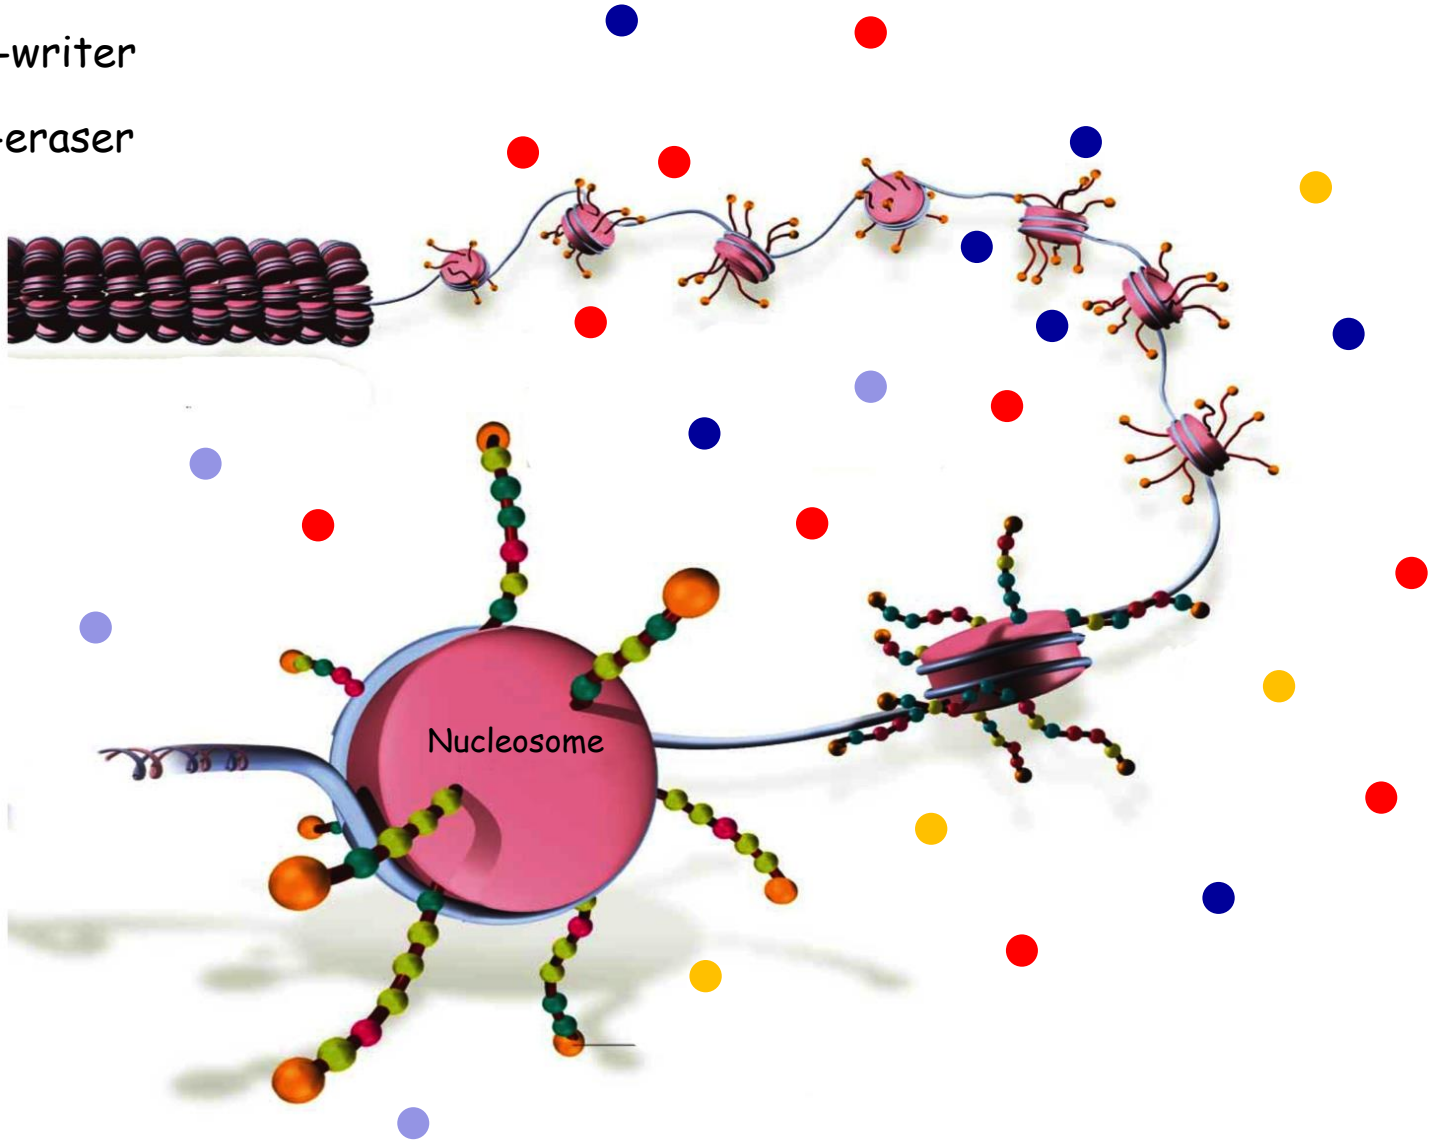

The epigenetic consensus problem

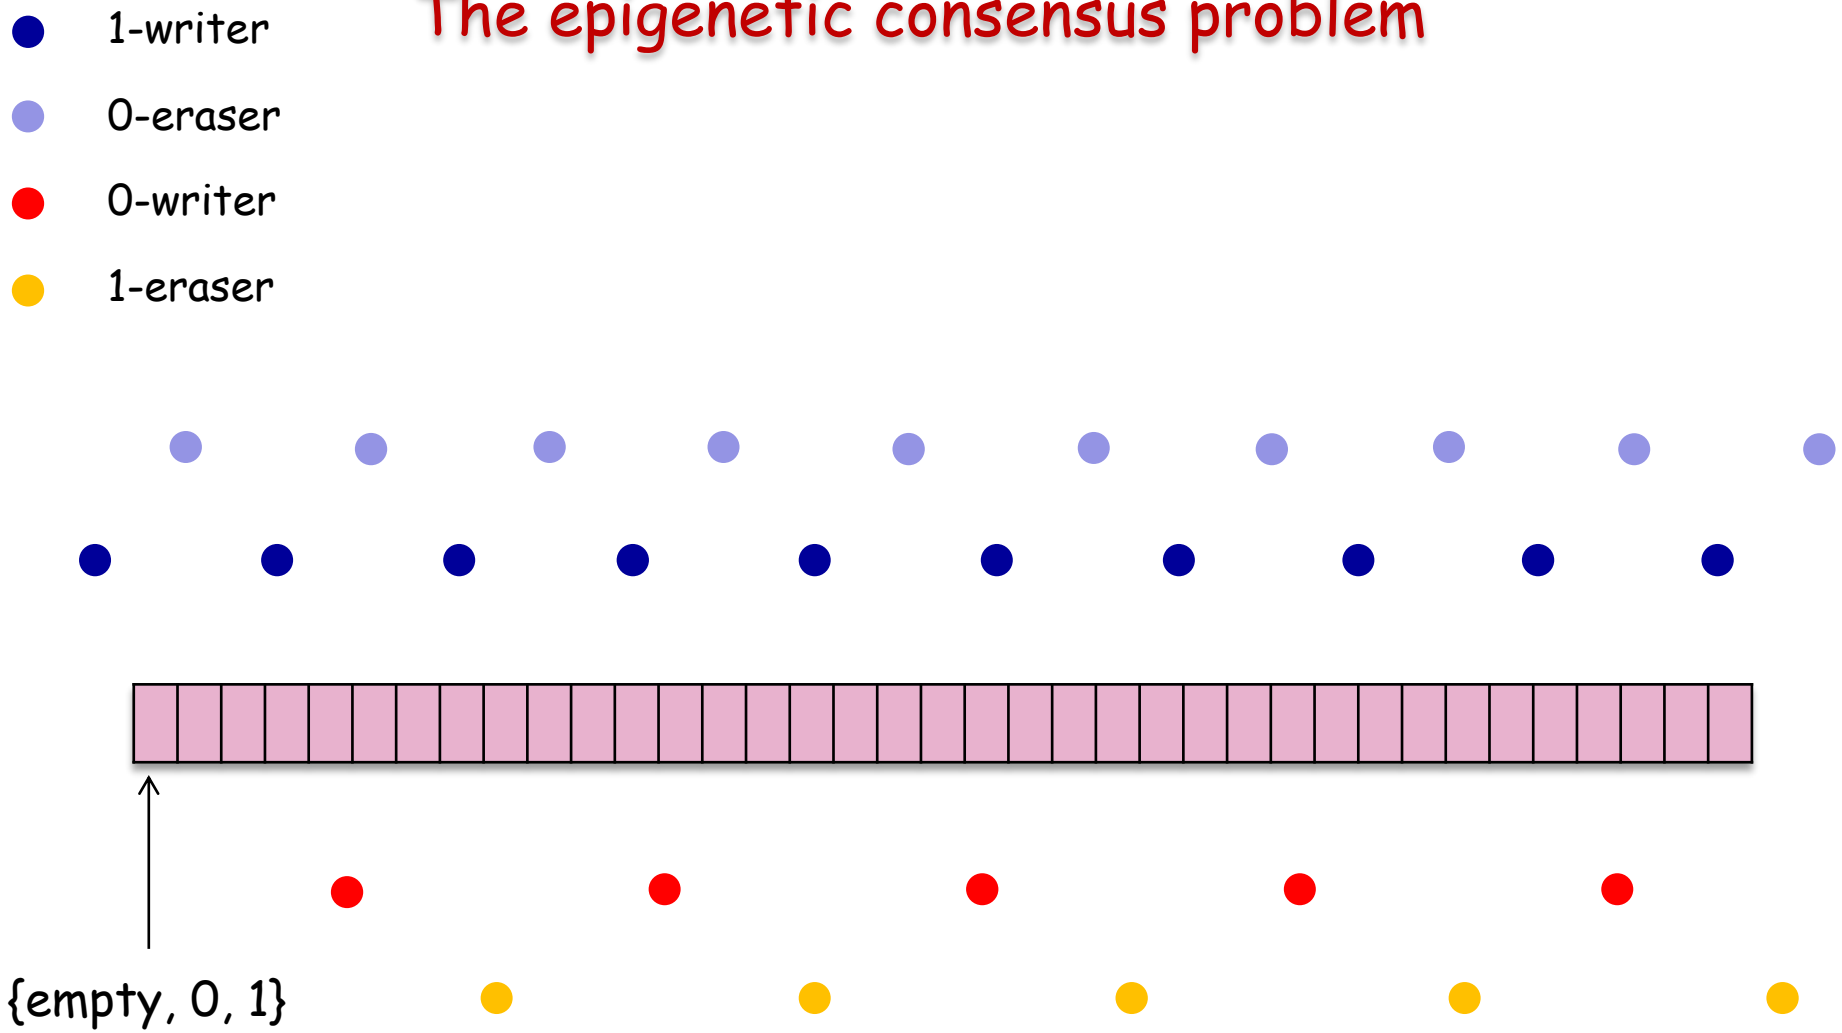

The epigenetic consensus problem

- 1-writer
- 0-eraser
- 0-writer
- 1-eraser

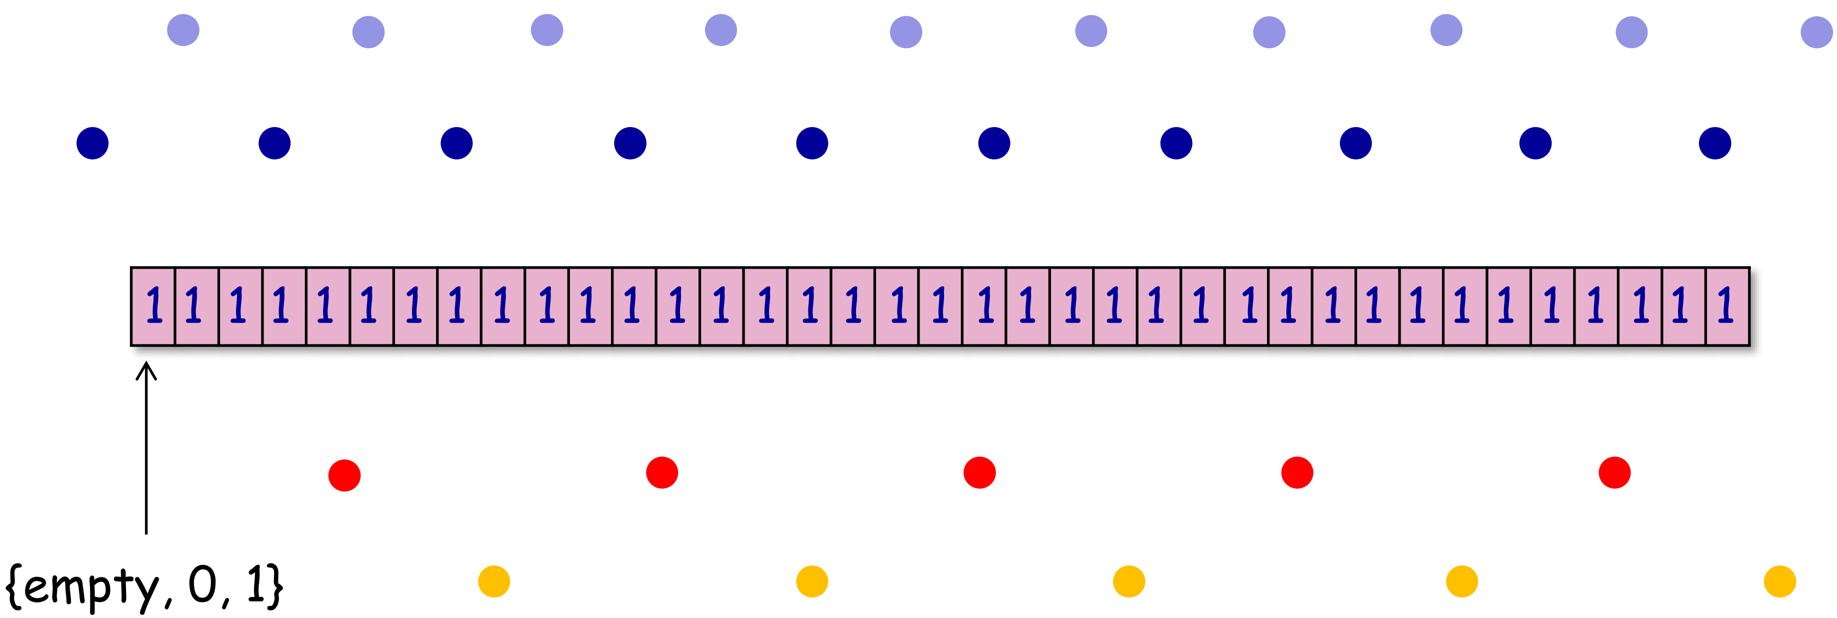

Very weak model

- Randomization
- Anonymous processes (no identifiers)
- Anonymous shared memory
- Memory-less processes (well may 1-2 bits)
- A transition from 0 to 1 cannot occur directly
- No sense of direction
- Self-stabilization

★ ★ ★ ★ ★

We present an algorithm that matches the biological assumptions, prove it correctness and derive bounds on its expected run time both theoretically and in simulations.

The epigenetic consensus problem

- 1-writer
- 0-eraser
- 0-writer
- 1-eraser

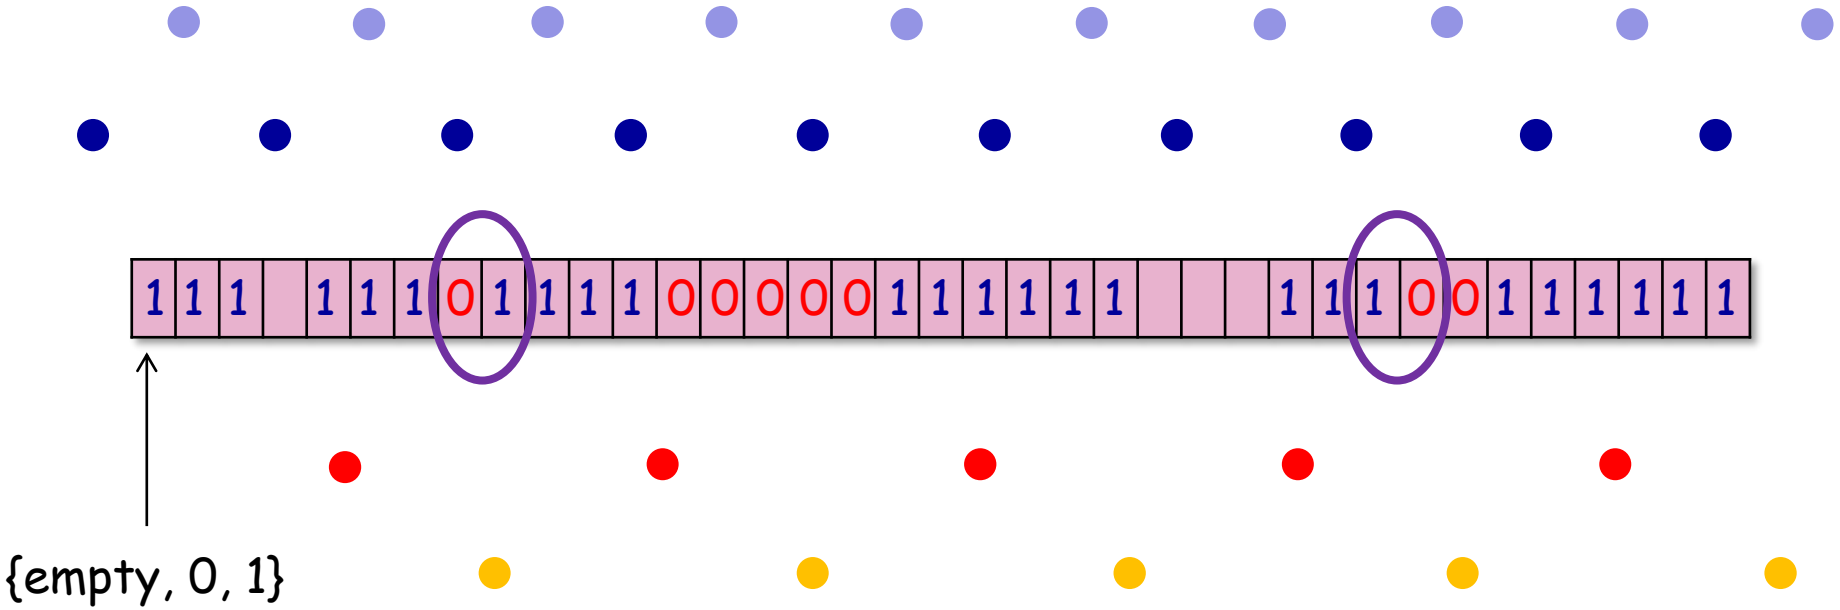

The epigenetic consensus problem

Corollary 6.2 *Assume $W_1/W_0 \geq 3$. The probability that the final decision value is 1 is more than*

$$\left(1-(1/3)^N\right) \times \left(1-e^{-N/12}\right)$$

Corollary 7.2 *Assuming $W_1/W_0 \geq 3$, $E[T] \leq 6.4N^2$*

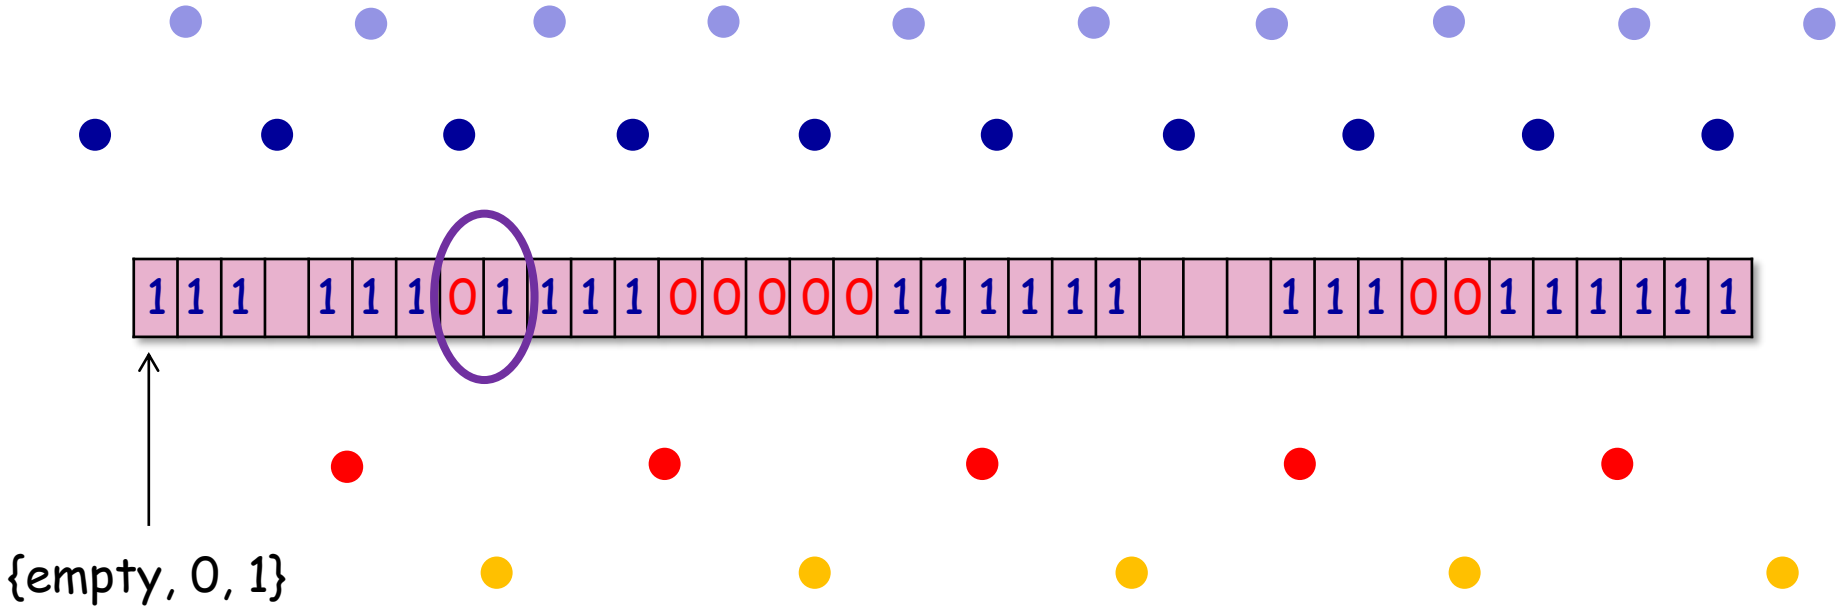

Annie's question: Where computers can help you except from text editor and a slides editor ?

Simulations

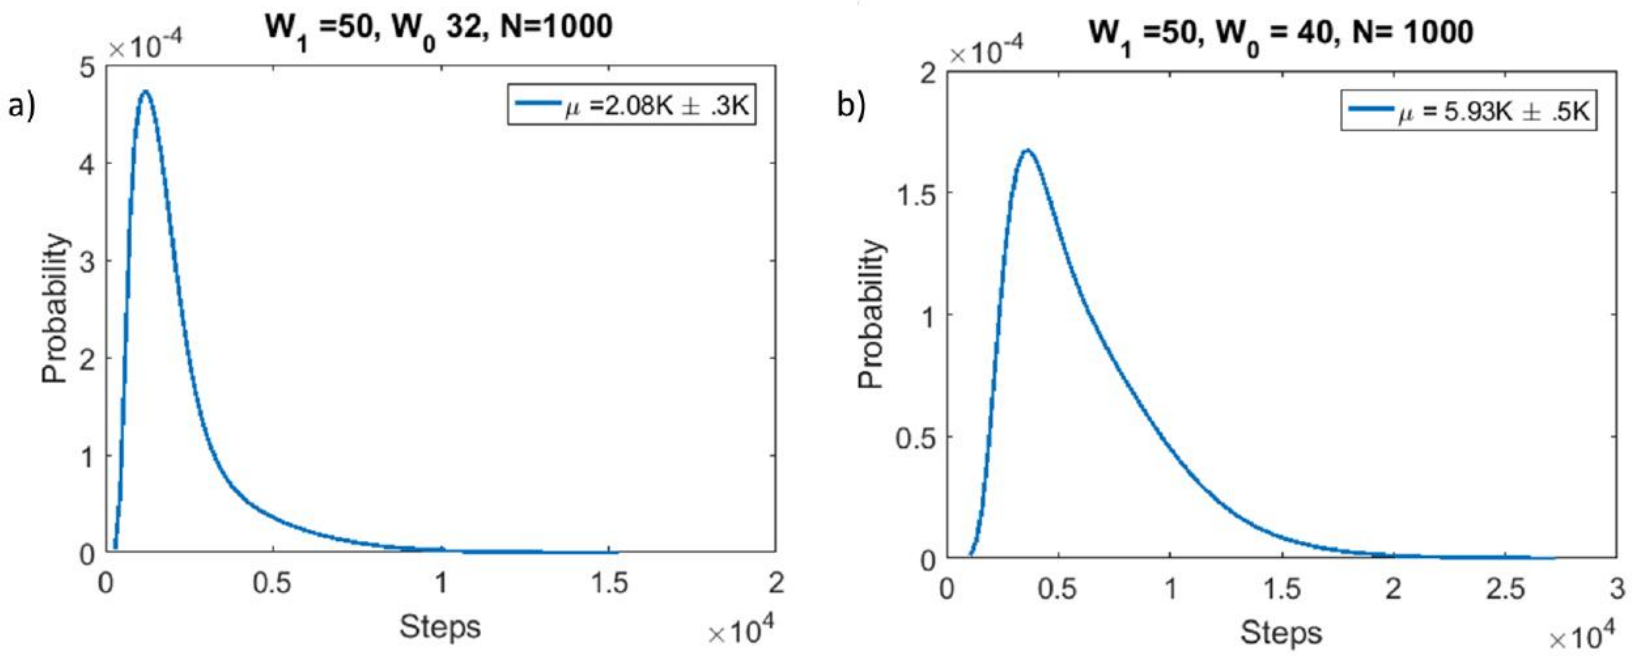

Figure 1: Distribution of number of steps to reach consensus. Plots summarize 300 random runs of the algorithm. a) low and b)high level of competitions between 1-writers and 0-writers. μ denotes the average time to reach consensus.

Simulations

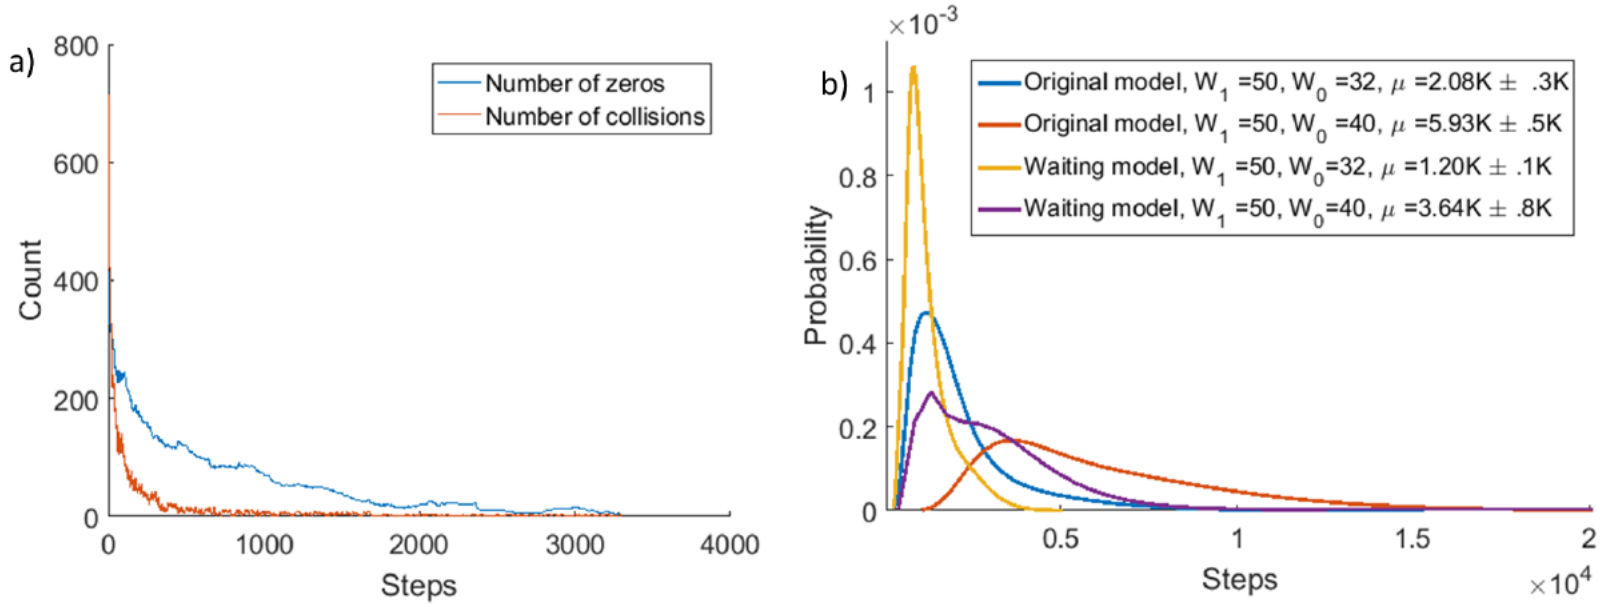

Figure 2: a) Number of zeros and collisions vs steps in the algorithm. While the initial number of collisions is a linear function of the number of 0's, we observe that towards the end of the algorithm there are very few collisions while the number of 0's remains relatively high. b) Comparison between the proposed original model and a revised model that allows writers to attach themselves with the erasers and the erasers wait until a collision is resolved. Here we can see that the waiting version is faster at both competition level compared to the original model.

Conclusion #1

Weak models are interesting!

Part II

Anonymous Shared Memory

Classical view of SM

Objects have names

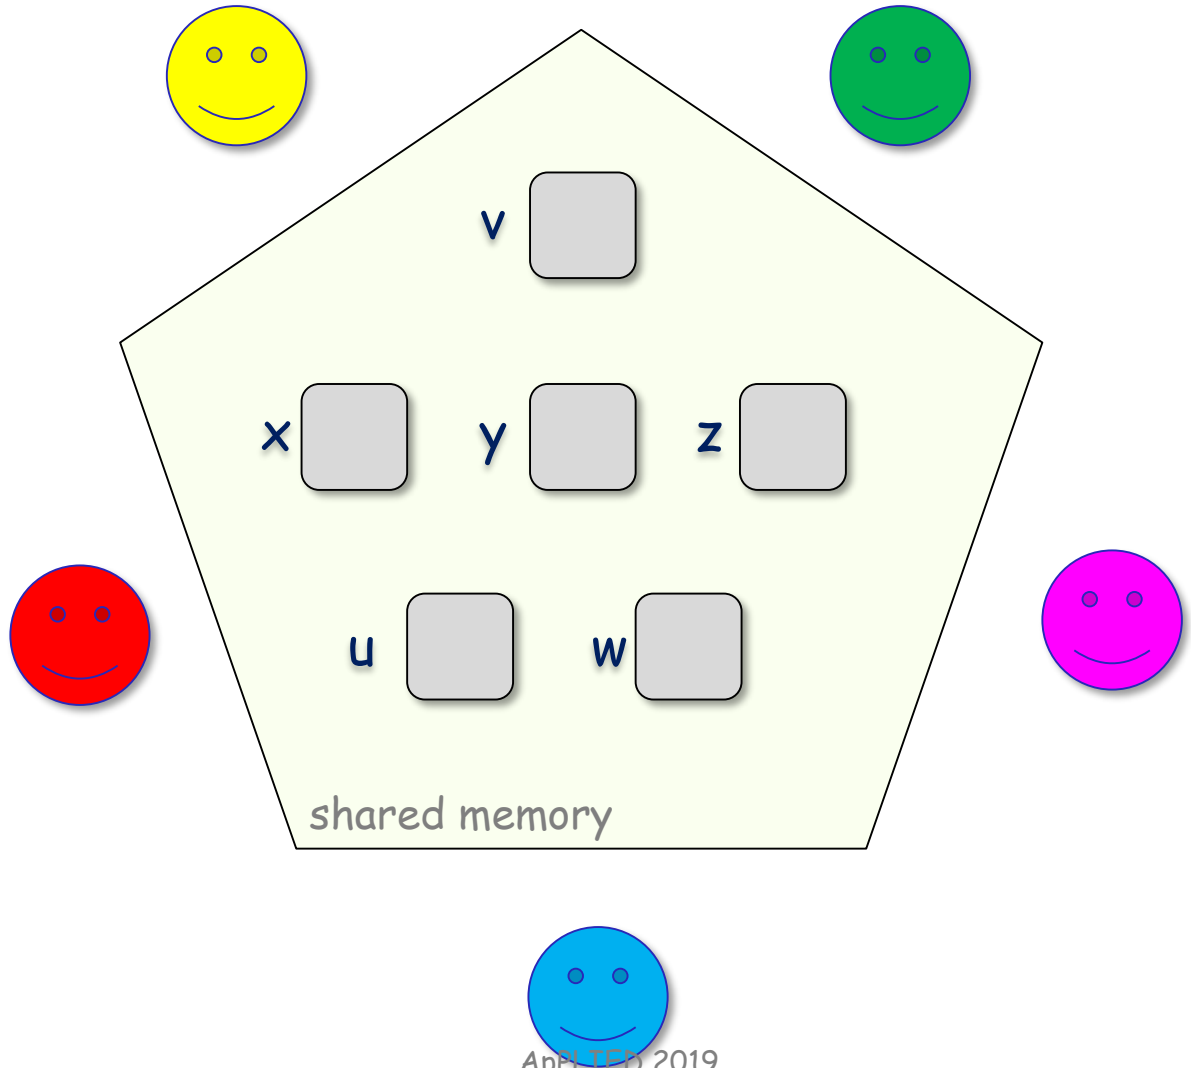

Anonymous shared memory

NO prior agreement on the names of the objects!

New Model

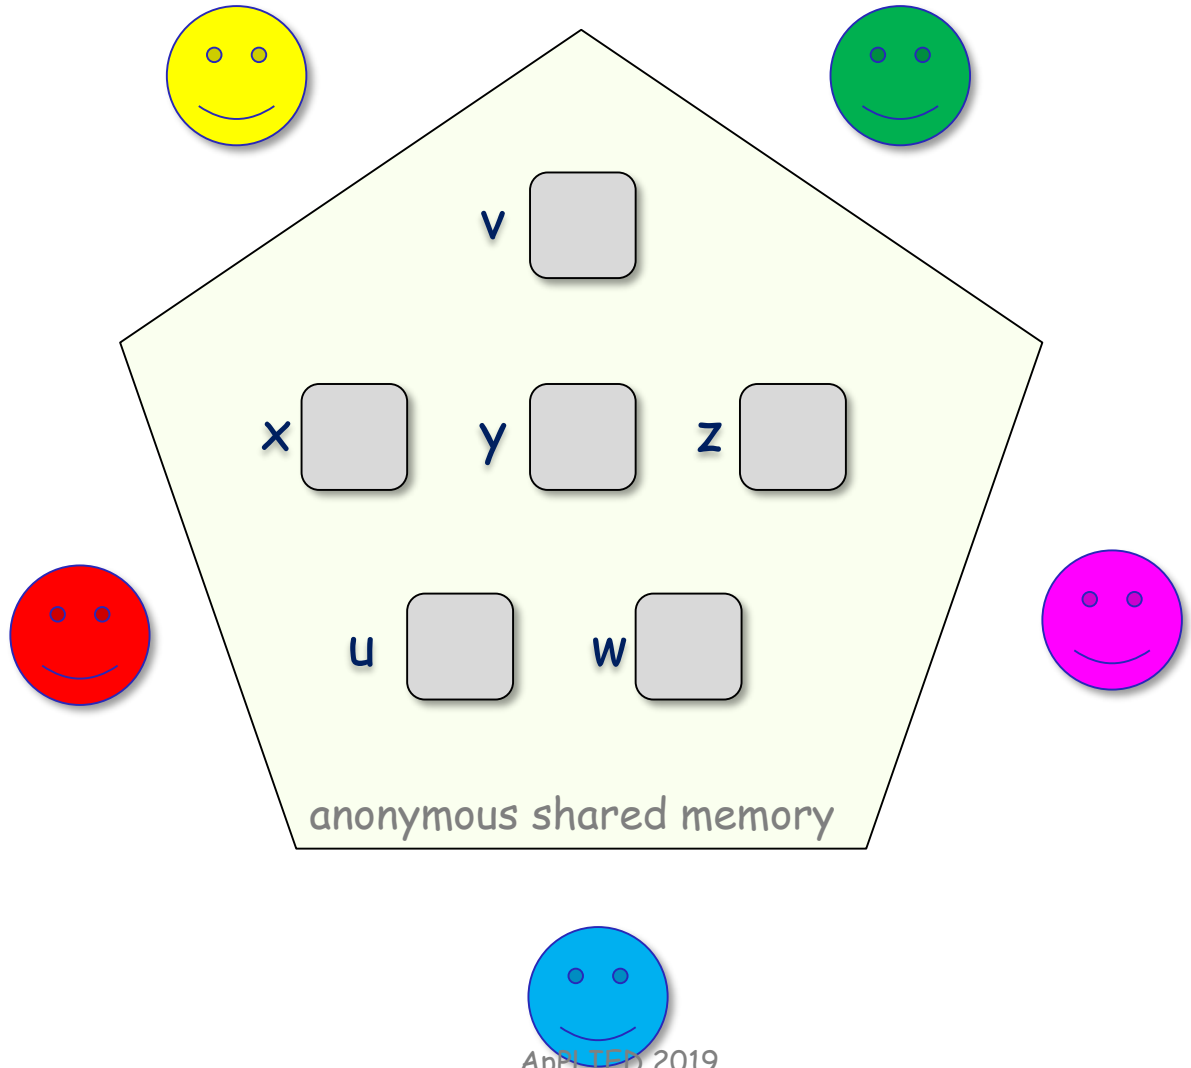

Anonymous shared memory

New Model

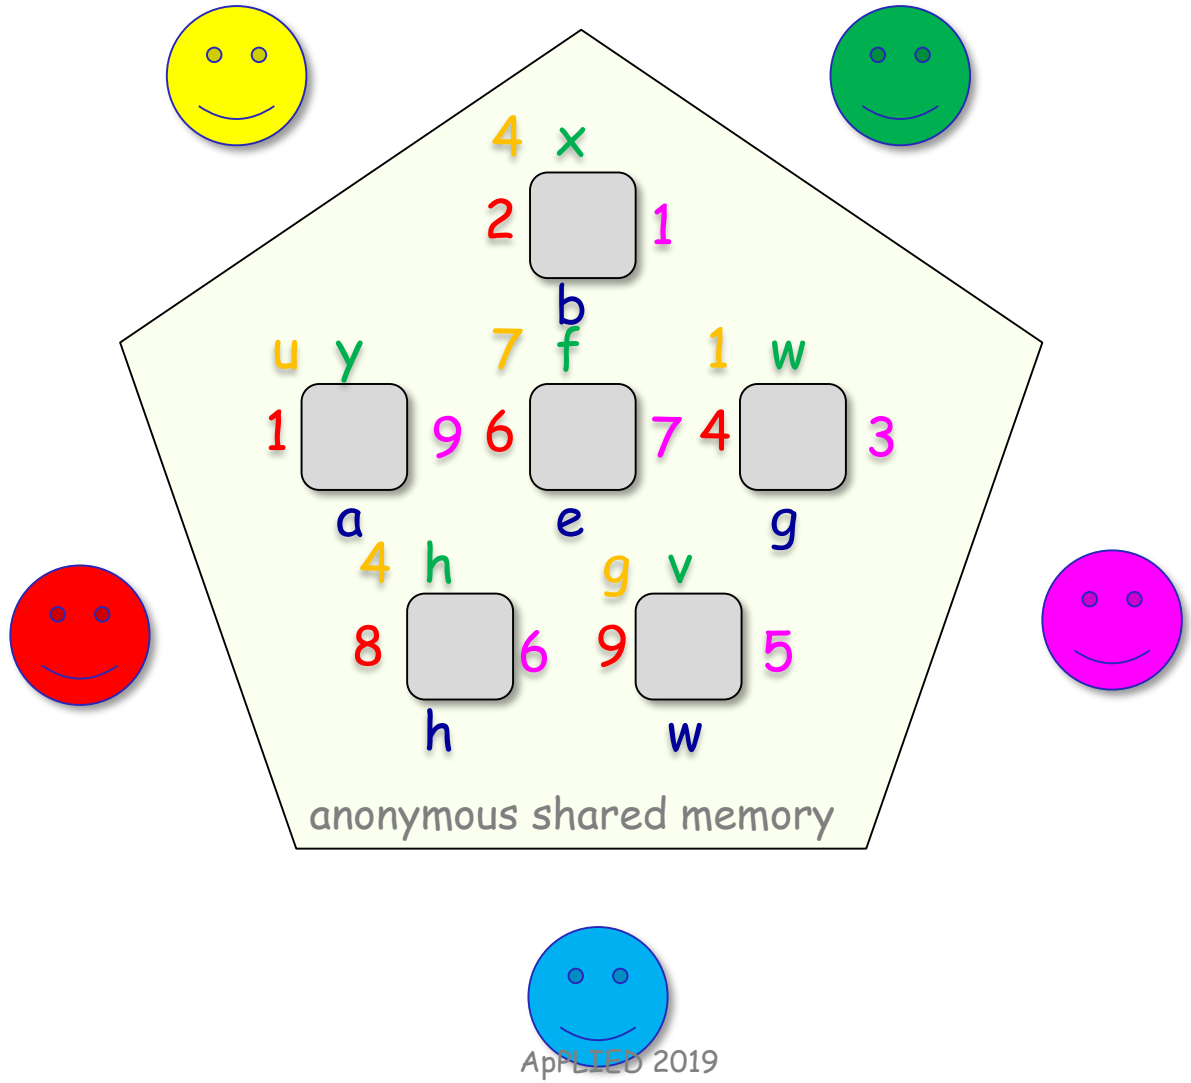

Anonymous shared memory

Coordination without prior agreement
by *Gadi Taubenfeld*

PODC 2017

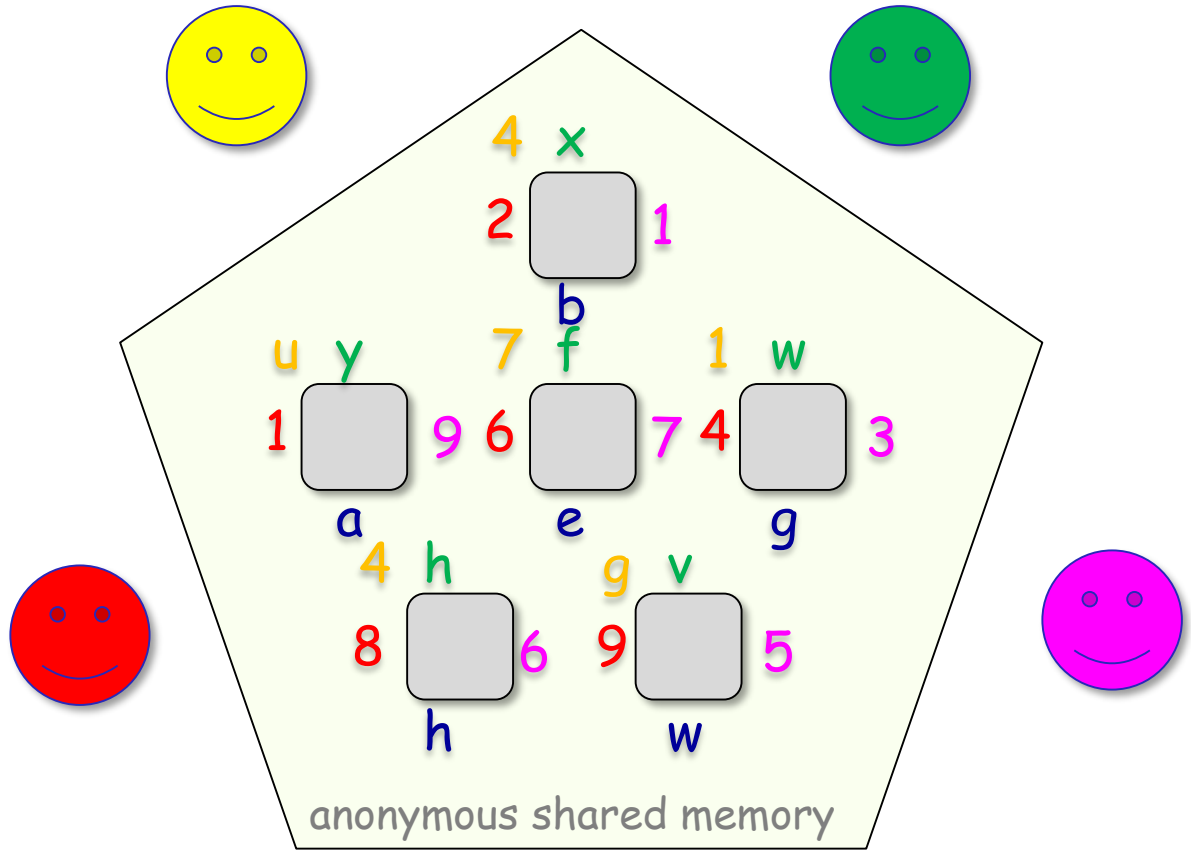

Algorithms & space bounds

PODC 2017
R/W registers

Algorithms	Can do ✓	Cannot do ✗
Deadlock-free symmetric mutual exclusion for two processes	odd # of registers	even # of registers
Obstruction-free consensus for n ≥ 2 processes	2n-1 or more	n or less
Obstruction-free adaptive perfect renaming for n ≥ 2 processes	2n-1 or more	n or less

(The # of registers is not 1)

PODC 2019

Optimal Memory-Anonymous Symmetric Deadlock-Free Mutual Exclusion

- Theorem. For every $n \geq 1$, there is a symmetric deadlock-free mutual exclusion algorithm for n processes using $m \geq 1$ anonymous R/W registers if and only if for every positive integer $1 < k \leq n$, m and k are relatively prime.

- The same result holds also for RMW registers ! *


Zahra Aghazadeh

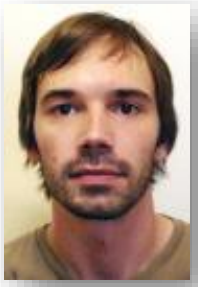
Damien Imbs


Michel Raynal


Philipp Woelfel


Gadi Taubenfeld

* It is trivial to do also with one RMW register.

Resolving two open problems

For a universe which includes (also) anonymous objects,

- Are atomic read/write registers the weakest objects ?
- Are deterministic (oblivious) objects with the same set-consensus number have the same computational power ?

NO

Conclusion #2

Weak models are interesting!

Part III

Fractions in Distributed Computing

Egypt
1600 B.C.

Europe
17th century

Dist. Comp.
???

Fractions were studied by Egyptians mathematicians around 1600 B.C.
However, fractions, as we use them today, didn't exist in Europe until the 17th century.

Part III

Fractions in Distributed Computing

- We understand what it means to tolerate one process failure.
- But what does it mean to tolerate 0.8 process failure ?

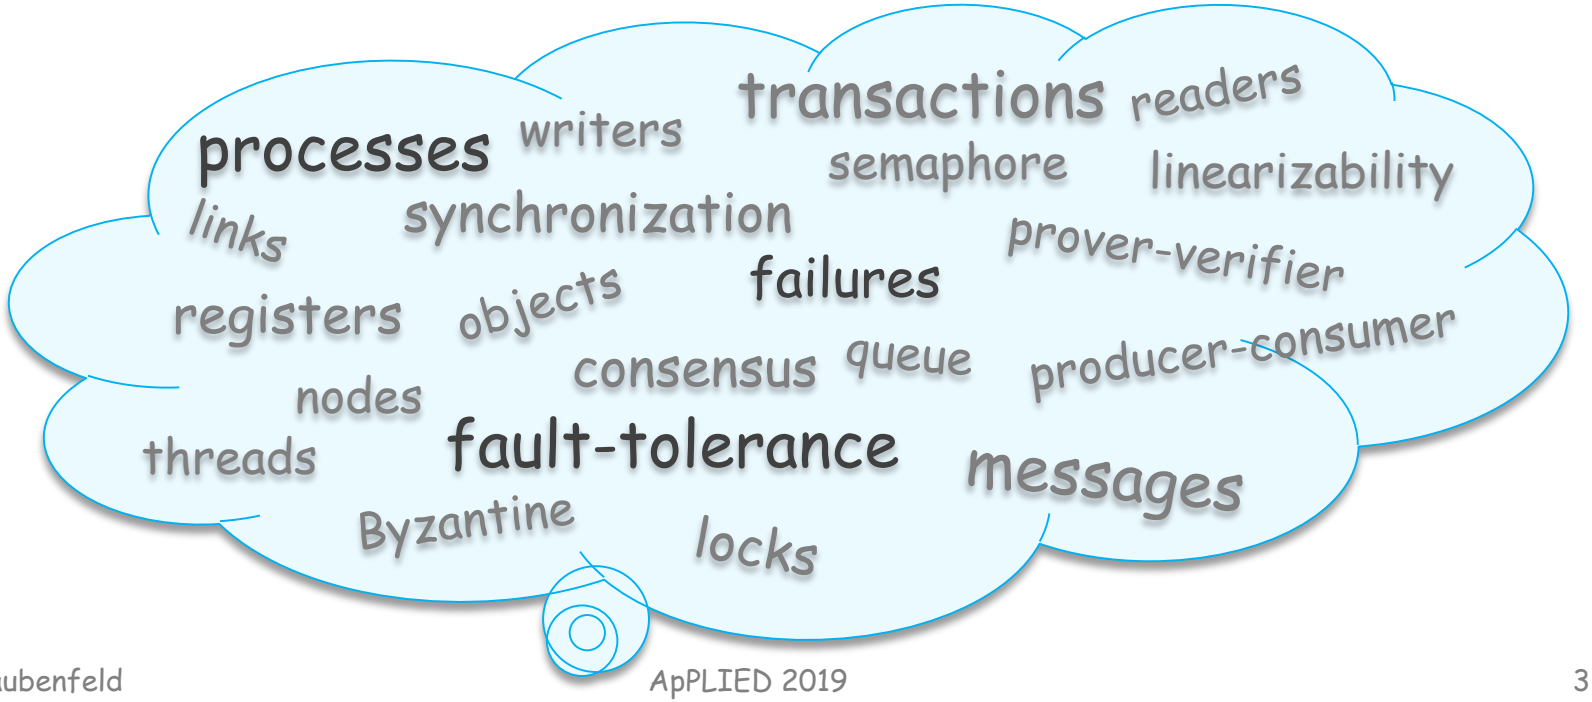

Motivation
Something is better than nothing

- FLP: Impossibility of consensus in the presence of a single failure.

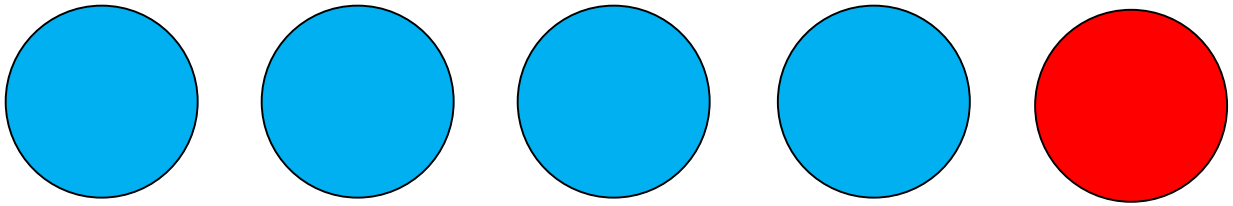

- Is consensus possible in the presence of a single weak failure?

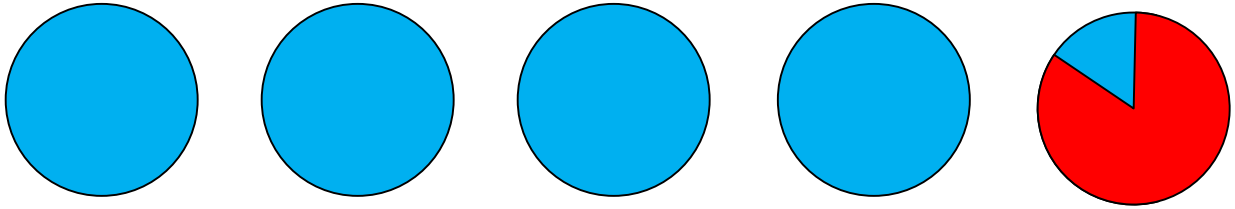

NETYS 2018

Weak Failures: Definitions, Algorithms and Impossibility Results by *Gadi Taubenfeld*

- Is consensus possible in the presence of a single <u>weak</u> failure?

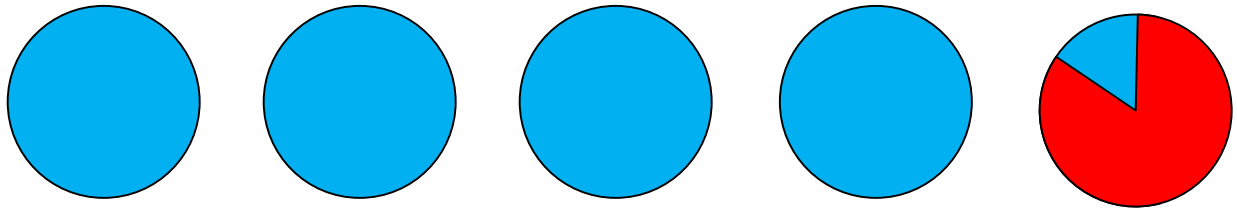

Motivation

Generalizing from the previous example

- Suppose you can solve a problem in the presence of f traditional failures, but not in the presence of f+1 such failures.

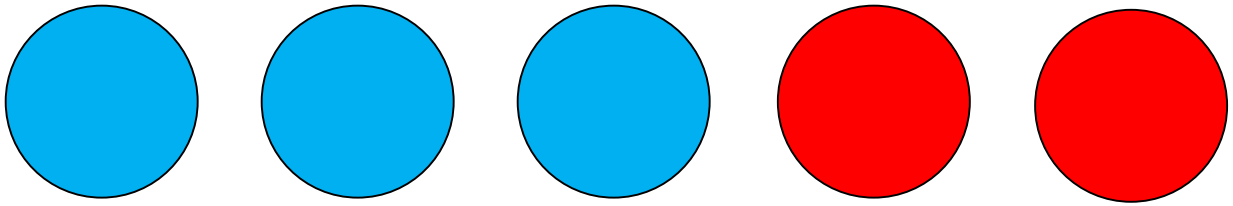

(f=2)

- Maybe it is possible to solve the problem in the presence of f traditional failures plus several weak failures.

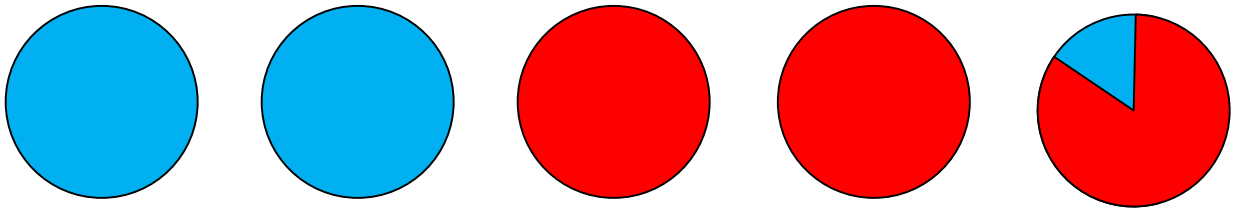

Set agreement and renaming in the presence of contention-related crash failures

SSS 2018

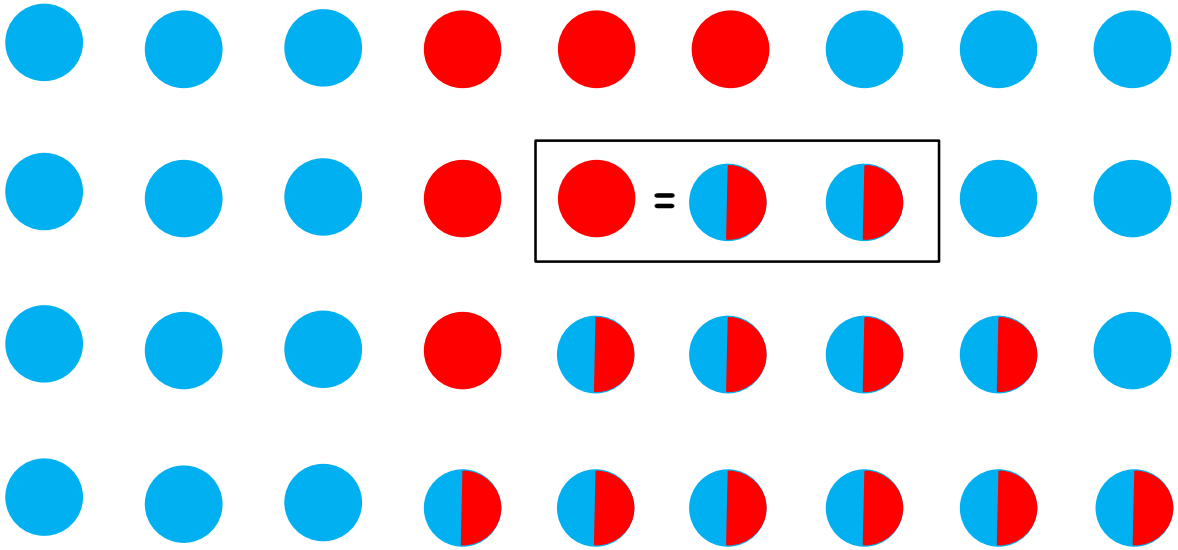

Anaïs Durand

Michel Raynal

Gadi Taubenfeld

Conclusion #3

Weak models are interesting!


THANKS