Natural Distributed Algorithms

- Lecture 1 -
Introduction to
(Noisy) Rumor Spreading
in (Desert) Ants

Emanuele Natale
CNRS - UCA

CdL in Informatica
Università degli Studi di Roma “Tor Vergata”

Ant Behavior is still Largely Unexplored

Ants invented **agricolture**

The ants cultivate species of fungi, feeding them with freshly cut plant material and keeping them free from pests and molds, and if they notice that a type of leaf is toxic to the fungus, they will no longer collect it. Some of this fungi no longer produce spores: they fully domesticated their fungal partner 15 million years ago.

Ants invented **war**

Some colonies conduct ritualized tournaments: Opposing colonies summon their worker forces to the tournament area, where hundreds of ants perform highly stereotyped display of fights. When one colony is considerably stronger than the other, the tournaments end quickly and the weaker colony is sacked. - The Ants, B. Hölldobler and E. O. Wilson

Ants invented **slavery**

Slave-making ants are brood parasites that capture broods of other ant species to increase the worker force of their colony. After emerging in the slave-maker nest, slave workers work as if they were in their own colony, while parasite workers only concentrate on replenishing the labor force from neighboring host nests.

Ant Behavior is still (Quite) Largely Unexplored

Ants invented **architectures**

When army ants need to cross a large gap, they simply build a bridge - with their own bodies. Linking together, the ants can move their living bridge from its original point, allowing them to cross gaps and create shortcuts across rainforests in Central and South America.

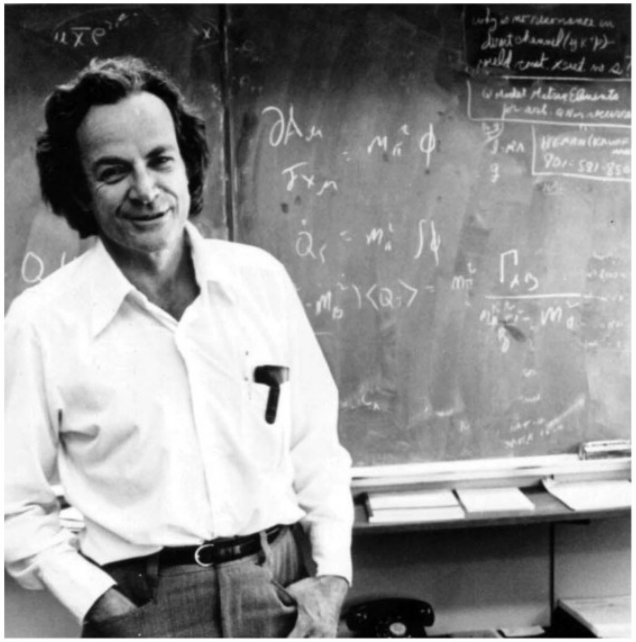

Ants puzzled **Feynman**

One question that I wondered about was why the ant trails look so straight and nice. The ants look as if they know what they're doing, as if they have a good sense of geometry. Yet the experiments that I did to try to demonstrate their sense of geometry didn't work. Many years later, when I was at Caltech . . .

And **more...**

Have a look at the many books (e.g. Hölldobler), or just Youtube.

Communication Noise in Biological Systems

"**Noise**" refers to a wide range of random or unwanted signals across different disciplines (acoustics, telecommunications, electronics, mathematics...)

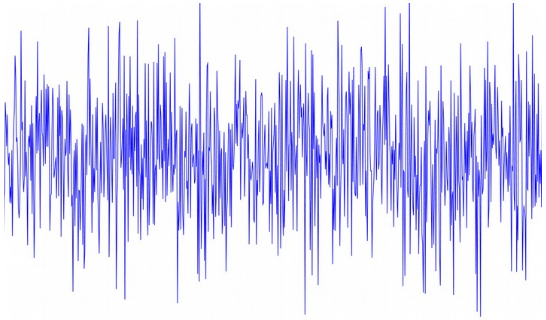

Information Theory
(Shannon 1948)
rigorously studies noise over communication channels

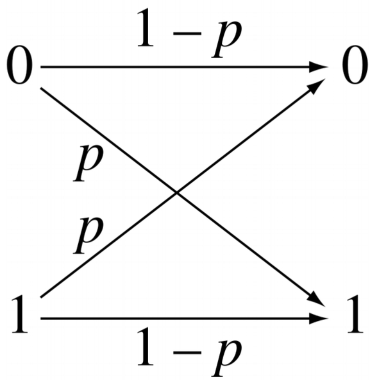

$$\mathcal{I}(X;Y) = H(Y) - H(p)$$

In **biology**, noise is widespread (cellular noise, developmental noise, neuronal noise...)

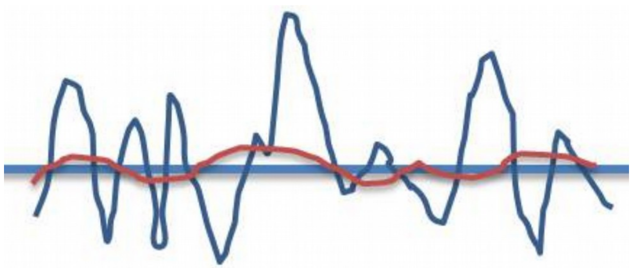

Neuronal Noise: administered signal and amplitude of noise

Motivation: Recruitment in Desert Ants

- **Cataglyphis niger** typical of Israel and other African countries
- Lives in desert (nest in dark cavities of rocks)
- Active during hottest hours
- Doesn’t use pheromones

Experimental setting

- A cricket is pinned near an artificial colony
- An ant finds it and needs to recruit nest mates to carry food.
- Data suggest that **communication is noisy** (reluctance to action).

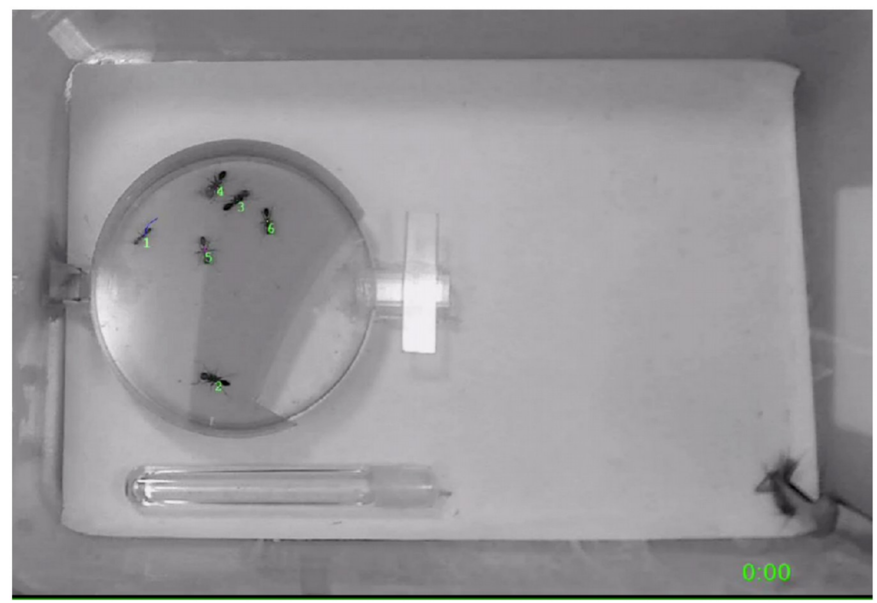

Video on catagliphys bombycina (saharan silver ant)

Experimental Setting

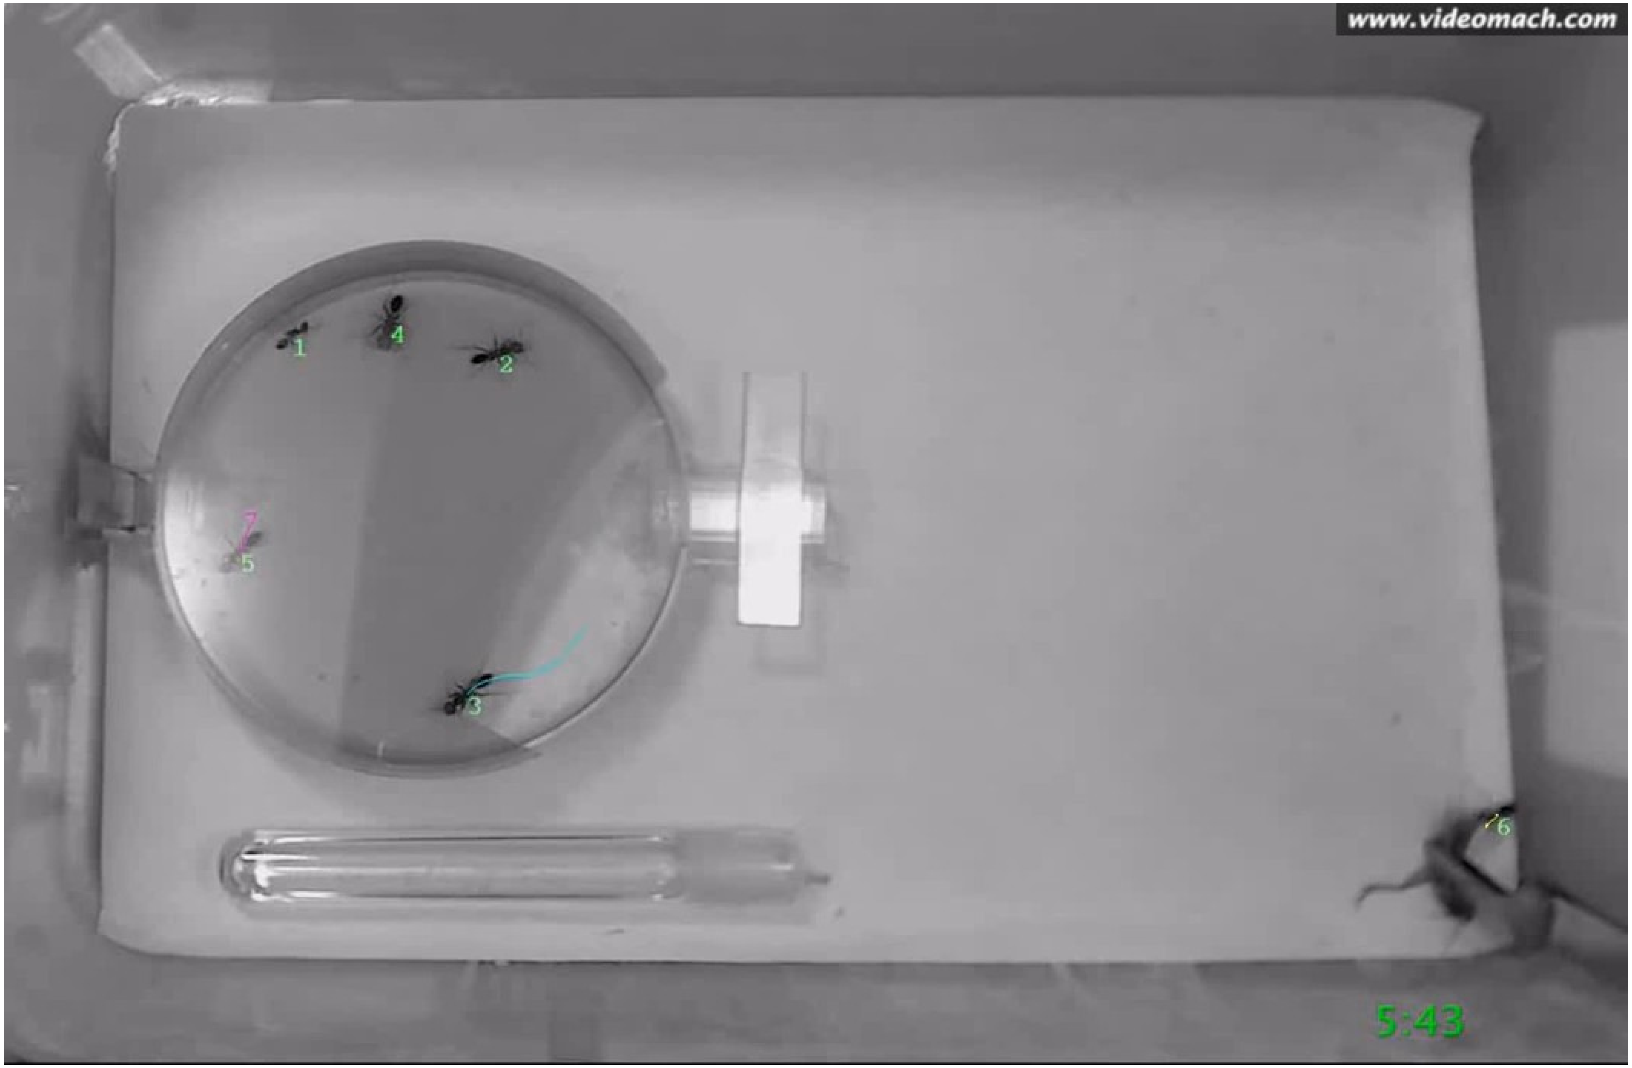

Stochastic Interactions: PUSH Model

Desired features
- Stochastic
- Parsimonious (Anonymous)
- Active (uni-directional)

(Uniform) PUSH model
[Demers ’88]
(single binary message)
- Discrete parallel time
- At each round each agent **can send one bit** message
- Each message is received by **one agent** chosen independently and uniformly at random

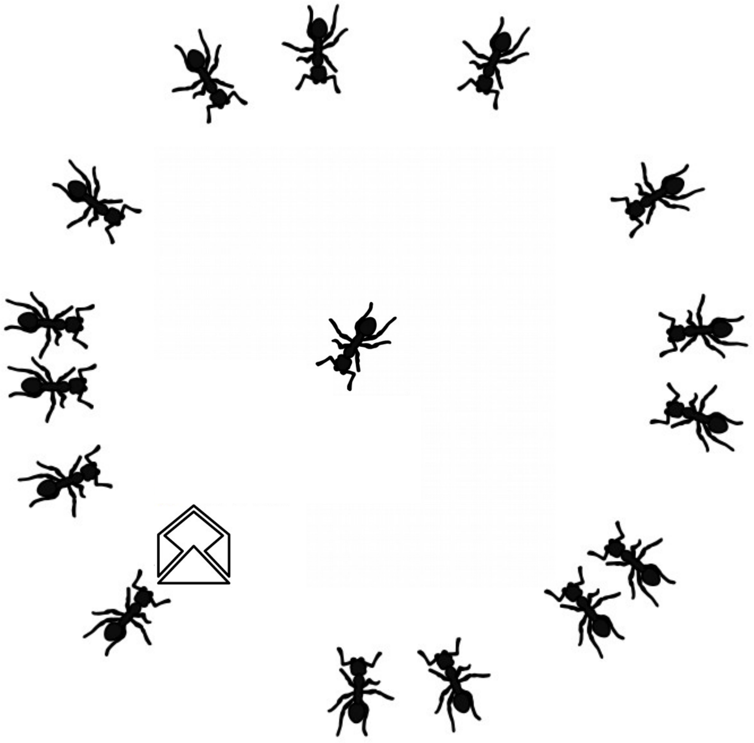

Stochastic Interactions: PULL Model

Desired features
- Stochastic
- Parsimonious (Anonymous)
- Passive (uni-directional)

(Uniform) PULL model
[Demers '88]
(single binary message)
- Discrete parallel time
- At each round each agent **shows one bit** message
- The agent can observe the message shown by **one agent** chosen independently and uniformly at random

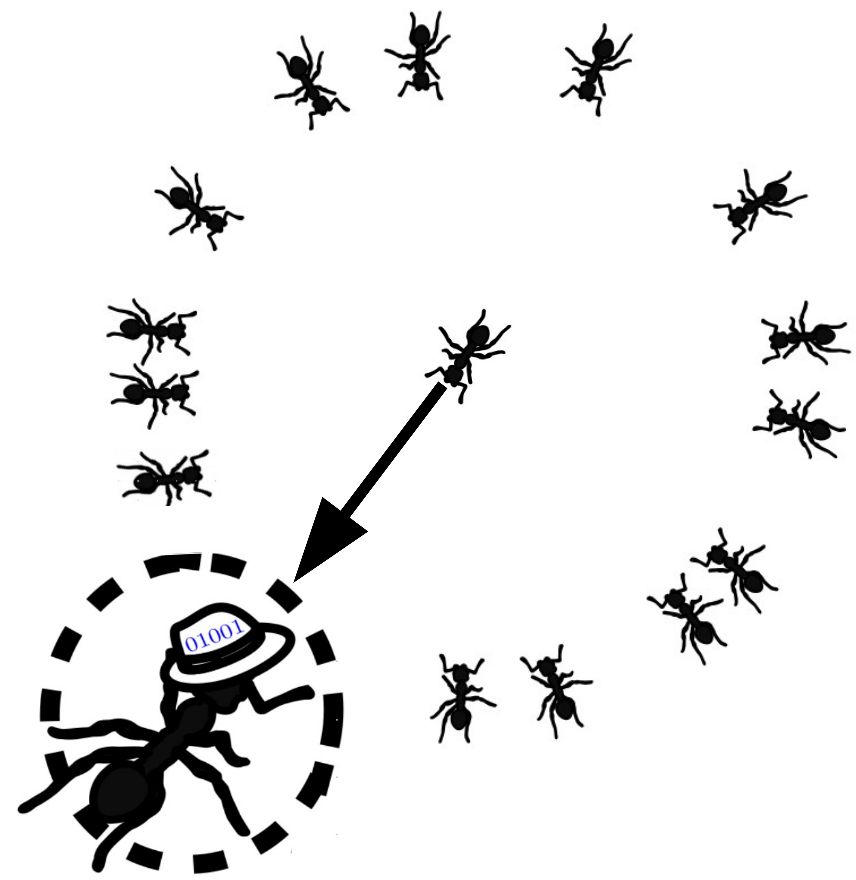

Noisy Communication

Before being received, each bit is flipped with probability $\frac{1}{2} - \epsilon$

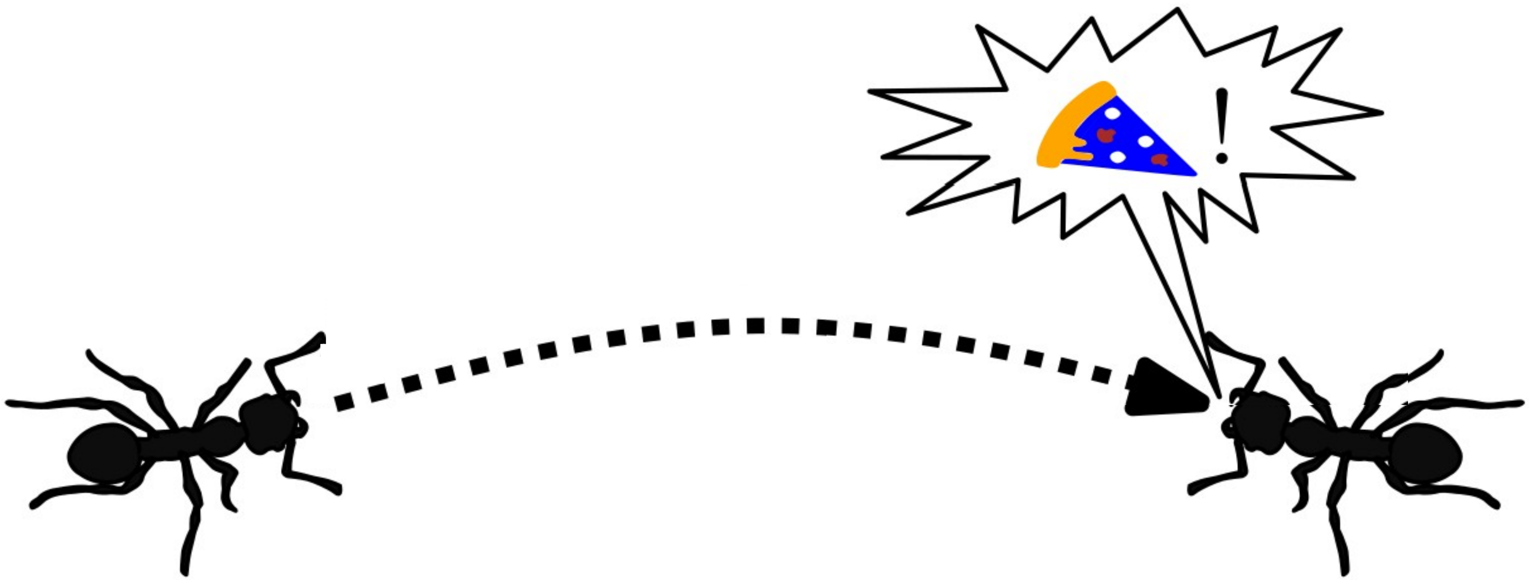

Rumor Spreading Problem

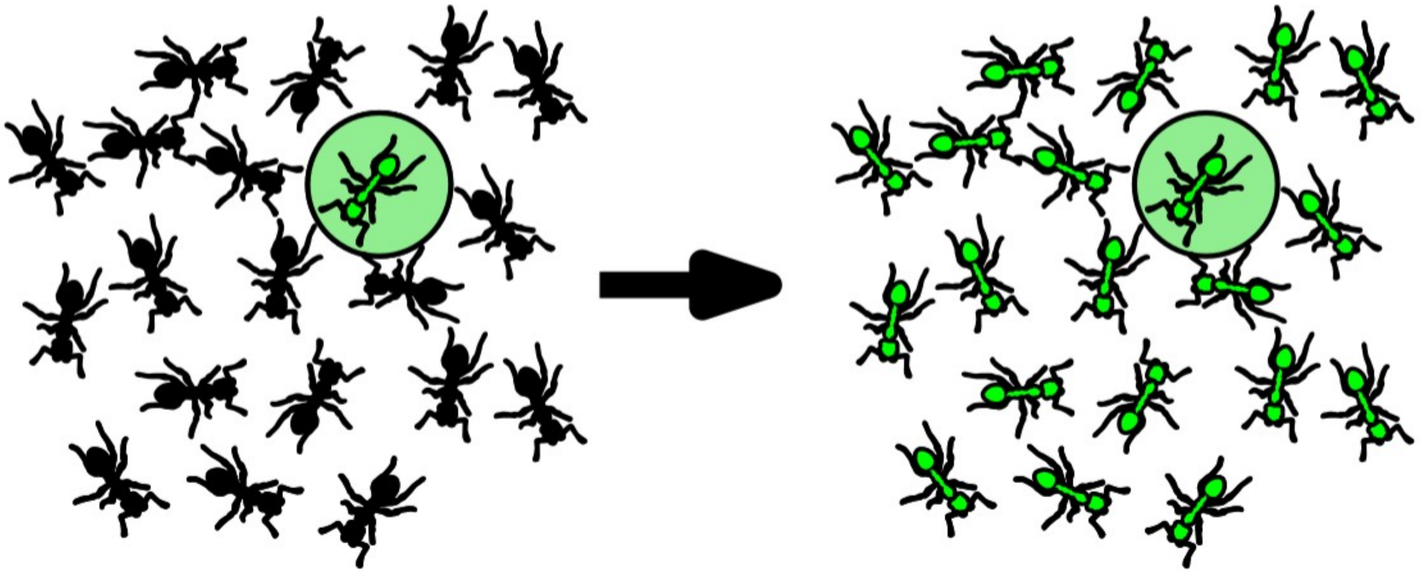

- One **source** node in a **special state**
- **Goal configuration**: all agents in the **special state**

Rumor Spreading in the (**Non-Noisy**) PUSH Model: "Chinese-Whispers" Protocol

How to solve the Rumor Spreading in the (Non-Noisy) PUSH Model?

Each agent sends the **rumor** as soon as it receives it

Rumor Spreading in the (Non-Noisy) PUSH Model: "Mean-Field Analysis"

How long does (Non-Noisy) Rumor Spreading take **in expectation** in the **PUSH** Model?

$\mathrm{I}^{(t)}$ number of **informed** nodes at time t. $I^{(0)} = 1$

$$\mathbb{E}\left[I^{(t)} \mid I^{(t-1)} = i\right] = i + (n-i)\left(1 - \left(\frac{n-1}{n}\right)^i\right)$$

already informed

not-yet informed

probability to become informed (at least one independent event)

When $I^{(t)}$ is small, what is a good lower bound?
(Hint: compute the expectation when $i = 1, 2, \ldots$

Problem. Show that the growth is **exponential** up to $\frac{i}{2}$

When $n - I^{(t)}$ is large, what is a good lower bound? (same hint...)

Problem. Show that the drop is **exponential** starting from $\frac{i}{2}$

How Good are Analyses in Expectation?

Expectation may be misleading:
Consider position of a symmetric random walk after t steps

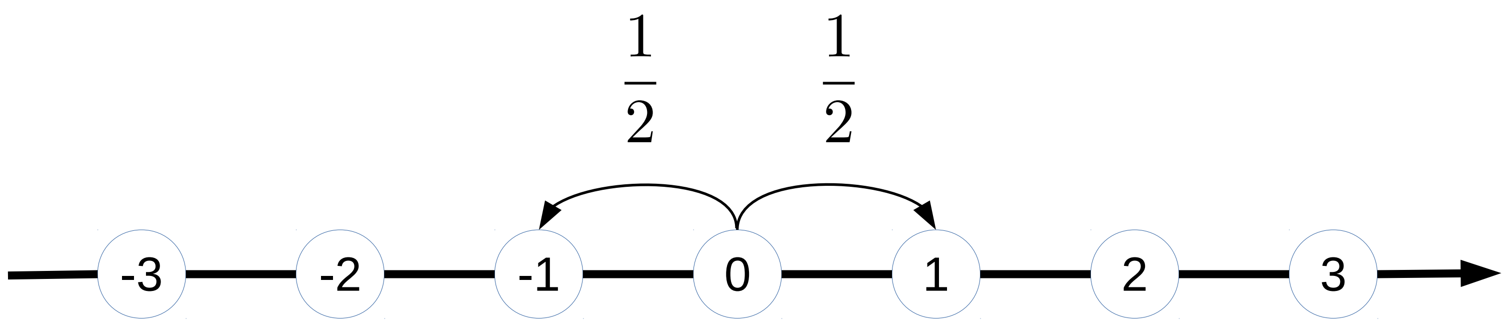

$$\mathbb{E}[\text{new position } X^{(t)} \mid \text{current position } X^{(t-1)} = x] = x$$

How long does Rumor Spreading take with high probability
in the PULL and PUSH Model?

Rumor Spreading in the (Non-Noisy) PUSH Model: Concentration of Probability?

General strategy: Apply Chernoff bounds, step-by-step, to show that the process follows w.h.p. its expected behavior.

(random variables)

(Un)Informed nodes can be modeled with binary r.v.s X_i

$$\Pr\left(\left|\sum_i X_i - \mathbb{E}\left[\sum_i X_i\right]\right| \geq \Delta\right) \leq 2e^{-\frac{\Delta^2}{2\left(\mathbb{E}[\sum_i X_i] + \frac{\Delta}{3}\right)}}$$

Second moment of binary r.v. equals expectation

Observe that, to obtain high probability, we need $\Delta \geq 2\sqrt{\mathbb{E}\left[\sum_i X_i\right] \log n}$ and at least $\mathbb{E}\left[\sum_i X_i\right] \geq 4 \log n$ but these “borderline” cases can be handled separately.

Problem. Unrolling the (w.h.p.-)recurrences still takes some work, but this can be part of a theory project.

Rumor Spreading in the (**Non-Noisy**) PULL Model: "Chinese-Whispers" Protocol

How to solve the Rumor Spreading in the (Non-Noisy) PULL Model?

Each agent asks for the **rumor** until it receives it

Rumor Spreading in the (Non-Noisy) PULL Model: "Mean-Field Analysis" (1/2)

How long does (Non-Noisy) Rumor Spreading take **in expectation** in the **PULL** Model? As for the PUSH, we have...

$\mathrm{I}^{(t)}$ number of **informed** nodes at time t. $\quad I^{(0)} = 1$

$$\mathbb{E}\left[I^{(t)} \mid I^{(t-1)} = i\right] = \underbrace{i}_{\text{already informed}} + \underbrace{(n-i)}_{\text{not-yet informed}} \underbrace{\frac{i}{n}}_{\text{probability to become informed}} = i\left(2 - \frac{i}{n}\right) \overset{\text{if } i \leq \frac{n}{2}}{\geq} i\frac{3}{2}$$

So by iterated expectation...

$$\mathbb{E}\left[I^{(t)}\right] = \mathbb{E}\left[\mathbb{E}\left[I^{(t)} \mid I^{(t-1)}\right]\right] \overset{\text{if } I^{(t-1)} \leq \frac{n}{2}}{\geq} \mathbb{E}\left[I^{(t-1)}\right]\frac{3}{2}$$

$$= \mathbb{E}\left[\mathbb{E}\left[I^{(t-1)} \mid I^{(t-2)}\right]\right]\frac{3}{2} \overset{\text{if } I^{(t-2)} \leq \frac{n}{2}}{\geq} \ldots \geq \left(\frac{3}{2}\right)^t$$

Rumor Spreading in the (Non-Noisy) PULL Model: “Mean-Field Analysis” (2/2)

We have proved $\mathbb{E}\left[I^{(O(\log n))}\right] \geq \frac{n}{2}$.

As for the number of uninformed agents

$$\mathbb{E}\left[U^{(t)} \mid U^{(t-1)} = u\right] = u - u\left(1 - \frac{u}{n}\right) = \frac{u^2}{n} \overset{\text{if } U^{(t-1)} \leq \frac{n}{2}}{\leq} \frac{u}{2}$$

Again, by the law of iterated expectation

$$\mathbb{E}\left[\mathbb{E}\left[U^{(t)} \mid U^{(t-1)}\right]\right] \overset{\text{if } U^{(t-1)} \leq \frac{n}{2}}{\leq} \frac{1}{2}\mathbb{E}\left[U^{(t-1)}\right]$$

$$\overset{\text{if } U^{(t-2)} \leq \frac{n}{2}}{\leq} \ldots \leq \frac{1}{2^t}$$

Rumor Spreading in the **Noisy** PUSH Model: "Chinese-Whispers" Does Not Work

How to solve the Rumor Spreading in the **Noisy** PUSH Model?

What if each agent sends the **rumor** as soon as it receives it?

Does “Greedy” (PUSH) Rumor Spreading Work?

PULL Model
(Stochastic
Interactions)

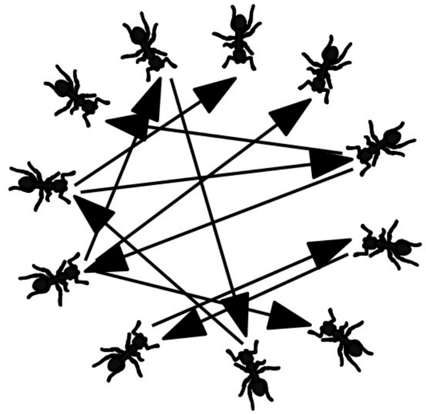

+

1 bit message
flipped
with prob 1/3
(Noisy
Communicaiton)

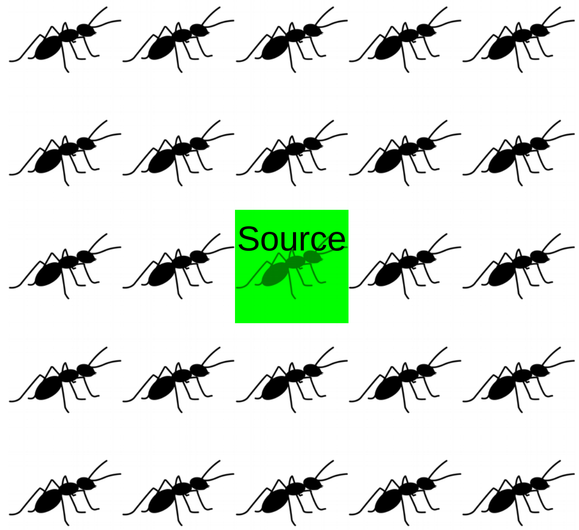

$$\frac{\color{red}{\text{red}}}{\color{green}{\text{green}}} = \frac{12}{25} = 0.48$$

The final **bias**
towards the right
opinion is close
to uniform!

"Greedy" (PUSH) Rumor Spreading Mixes Too Fast (1/2)

Let's make the previous intuition a bit more rigorous.

When an agent receives the information, what is the **distance of an agent from the source?**

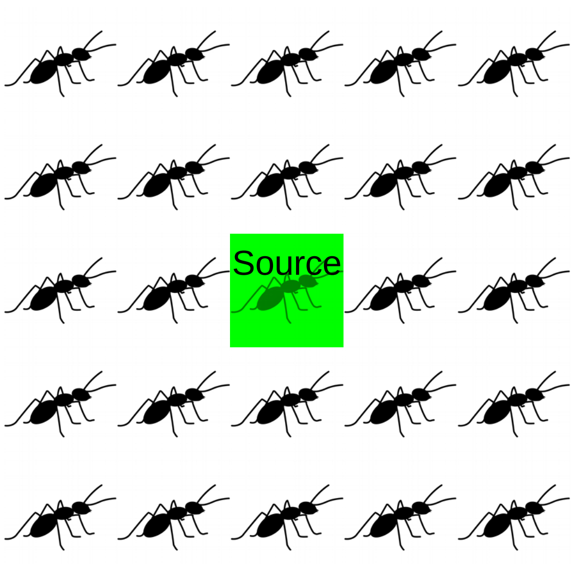

The growth is exponential until $o(n)$ agents are informed:

$$I^{(t)} \approx (1 + \text{positive constant})^t$$

approximately

The majority of agents receive the message after $\Omega\left(\log n\right)$ steps

“Greedy” (PUSH) Rumor Spreading Mixes Too Fast (2/2)

If $M^{(t)}$ is the message sent at time t

$$\Pr\left(M^{(t)} = 1\right) = \frac{2}{3}\Pr\left(M^{(t-1)} = 1\right) + \frac{1}{3}\Pr\left(M^{(t-1)} = 0\right)$$
$$\Pr\left(M^{(t)} = 0\right) = \frac{1}{3}\Pr\left(M^{(t-1)} = 1\right) + \frac{2}{3}\Pr\left(M^{(t-1)} = 0\right)$$

that is, defining $p_t := \Pr\left(M^{(t)} = 1\right)$
and $q_t := 1 - p_t := \Pr\left(M^{(t)} = 0\right)$

$$\begin{pmatrix} p_t \\ q_t \end{pmatrix} = \begin{pmatrix} \frac{2}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix} \begin{pmatrix} p_{t-1} \\ q_{t-1} \end{pmatrix}$$
$$= \begin{pmatrix} \frac{2}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{2}{3} \end{pmatrix}^t \begin{pmatrix} p_0 \\ q_0 \end{pmatrix} \approx \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

Project Idea 1

Write an overview on *Ant Algorithms* based on this course and

- Bruckstein, Alfred M. 1993. “Why the Ant Trails Look so Straight and Nice.” The Mathematical Intelligencer 15 (2): 59–62. https://doi.org/10.1007/BF03024195.
- Boczkowski, Lucas, Emanuele Natale, Ofer Feinerman, and Amos Korman. 2018. “Limits on Reliable Information Flows through Stochastic Populations.” PLOS Computational Biology 14 (6): e1006195. https://doi.org/10.1371/journal.pcbi.1006195.
- Feinerman, Ofer, Bernhard Haeupler, and Amos Korman. 2017. “Breathe before Speaking: Efficient Information Dissemination despite Noisy, Limited and Anonymous Communication.” Distributed Computing 30 (5): 339–55. https://doi.org/10.1007/s00446-015-0249-4.
- Fonio, Ehud, Yael Heyman, Lucas Boczkowski, Aviram Gelblum, Adrian Kosowski, Amos Korman, and Ofer Feinerman. 2016. “A Locally-Blazed Ant Trail Achieves Efficient Collective Navigation despite Limited Information.” ELife 5 (November): e20185. https://doi.org/10.7554/eLife.20185.
- Emek, Yuval, Tobias Langner, Jara Uitto, and Roger Wattenhofer. 2014. “Solving the ANTS Problem with Asynchronous Finite State Machines.” In ICALP, 471–82. http://link.springer.com/chapter/10.1007/978-3-662-43951-7_40.

Hints on difficulty: little or no math to deal with but lots to read and write.

Project Idea 2

Learn the complete analysis of the non-noisy rumor spreading in the PUSH model. In particular:

- These slides contain a lot of details that may have not been covered during the lecture (check the hyperlinks!)
- We didn't deal with the case $I^{(t-1)} = O\left(\log n\right)$...
  - Hint 1: How many nodes are informed **directly from the source agent** during the first $8 \log n$ rounds?
- When at least half of the agents are informed, what is the probability that an uniformed agent remains uninformed for $\log n$ rounds? After answering, apply the union bound.
- In order to unroll the recurrence relations of the bound on informed agents w.h.p., it is necessary to apply the **chain rule** and observe that **the number of informed agent can at most double**.
- **Warning:** There is a subtle problem in applying the Chernoff bound in the PUSH model. You can ignore it but you realize where the problem is.

Clip from BCC Silver Ant

Go back

“w.h.p.”

We say that an event holds with high probability if it holds with probability $1 - n^{-\Theta(1)}$.

Main tool to prove high probability: Chernoff bounds **(CBs)**.

Let X_i be independent r.v.s with $|X_i| \leq M$, then

$$\Pr\left(\left|\sum_i X_i - \mathbb{E}\left[\sum_i X_i\right]\right| \geq \Delta\right) \leq 2e^{-\frac{\Delta^2}{2\left(\sum_i \mathbb{E}[X_i^2] + \frac{M}{3}\Delta\right)}}$$

Proof. Chung, Fan, and Linyuan Lu. 2006. “Concentration Inequalities and Martingale Inequalities: A Survey.” Internet Mathematics 3 (1): 79–127. https://projecteuclid.org/euclid.im/1175266369.

Example:

Let $X_1, \ldots, X_n$ s.t. $\left(\Pr\left(X_i = 1\right) = \Pr\left(X_i = 0\right) = \frac{1}{2}\right)$ then

$$\Pr\left(\left|\sum_i X_i - \frac{n}{2}\right| \geq \sqrt{n \log n}\right) \leq 2e^{-\frac{n \log n}{\frac{n}{2} + \frac{1}{3}\sqrt{n \log n}}} \ll e^{-\log n} = \frac{1}{n}$$

Chernoff Bounds (CBs)

To prove CB, use **Markov's inequality**
Given positive random variable X

$$\mathbb{E}\left[X\right] = \sum_{i=1}^{\infty} i \Pr\left(X = i\right) \geq \sum_{i=t}^{\infty} i \Pr\left(X = i\right)$$

$$\geq \sum_{i=t}^{\infty} t \Pr\left(X = i\right) = t \Pr\left(X \geq t\right)$$

Use $a \leq b \iff e^a \leq e^b$ to make the random variable positive

For any Y we have

$$\Pr\left(Y \geq t\right) = \Pr\left(e^Y \geq e^t\right) \leq \frac{\mathbb{E}\left[e^Y\right]}{e^t}$$

Introduce a positive parameter λ and optimize

$$\Pr\left(Y \geq t\right) = \Pr\left(\lambda Y \geq \lambda t\right) \leq \inf_{\lambda} \frac{\mathbb{E}\left[e^{\lambda Y}\right]}{e^{\lambda t}}$$

Probability of at-Least-One of i.i.d. Events

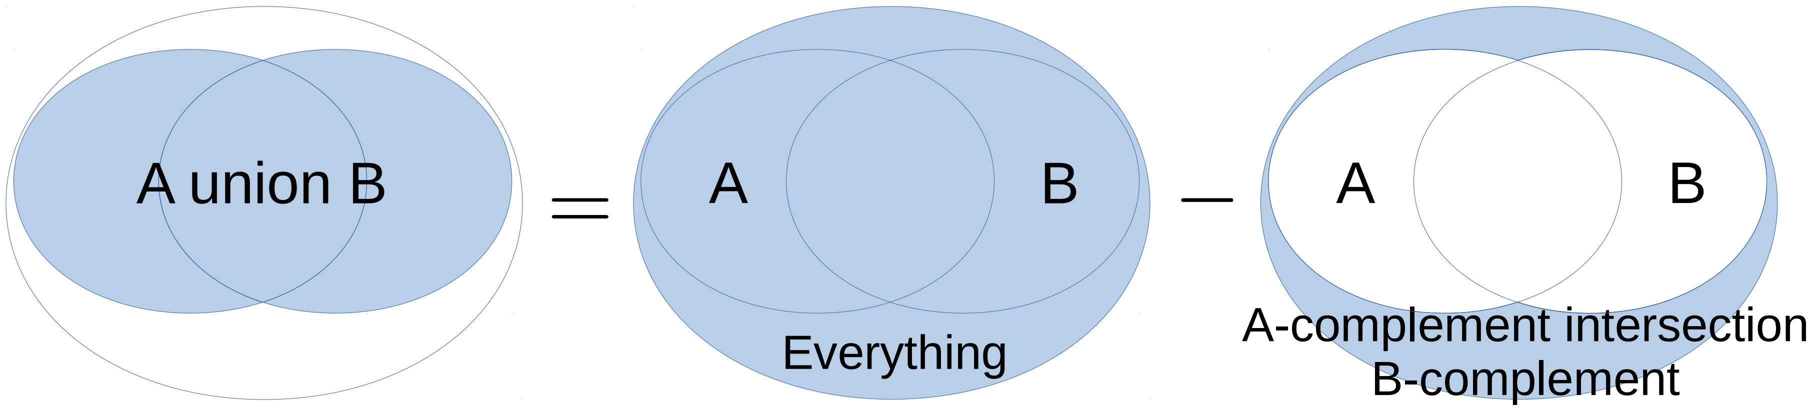

$$\Pr\left(\text{“}X_1 = 1\text{”} \vee \ldots \vee \text{“}X_n = 1\text{”}\right)$$

(complement event) $= 1 - \Pr\left(\text{“}X_1 = 0\text{”} \wedge \ldots \wedge \text{“}X_n = 0\text{”}\right)$

(by independence) $= 1 - \Pr\left(\text{“}X_1 = 0\text{”}\right) \wedge \ldots \wedge \Pr\left(\text{“}X_n = 0\text{”}\right)$

(same distribution) $= 1 - \Pr\left(\text{“}X_1 = 0\text{”}\right)^n$

Lower Bound on Informed Agents (1/2)

Using $1-y \leq e^{-y}$ we have $1-\left(1-\frac{1}{n}\right)^i \geq 1-e^{-\frac{i}{n}}$, that is

$$\mathbb{E}\left[I^{(t)} \,|\, I^{(t-1)} = i\right] \geq i + (n-i)\left(1-e^{-\frac{i}{n}}\right),$$

and we can write the r.h.s. as $i\left(1+\left(\frac{n}{i}-1\right)\left(1-e^{-\frac{i}{n}}\right)\right)$.

Computations show that

$$\frac{d}{di}\left(\frac{n}{i}-1\right)\left(1-e^{-\frac{i}{n}}\right) \geq 0 \iff \left(\frac{n}{i}-1\right)-\frac{n^2}{i^2}\left(e^{\frac{i}{n}}-1\right) \geq 0$$
$$\iff -\frac{i^2}{n^2}+\frac{i}{n}-e^{\frac{i}{n}}+1 \geq 0$$

and the second derivative is

$$\frac{d}{di}\left(-\frac{i^2}{n^2}+\frac{i}{n}-e^{\frac{i}{n}}+1\right) = -2\frac{i}{n}+1-e^{\frac{i}{n}}$$

Back to PUSH

Lower Bound on Informed Agents (2/2)

Since $\frac{d}{di}(-\frac{i^2}{n^2}+\frac{i}{n}-e^{\frac{i}{n}}+1) \leq 0$ and $-\frac{i^2}{n^2}+\frac{i}{n}-e^{\frac{i}{n}}+1=0$ for $i=0$, then

$$\frac{d}{di}\left(\frac{n}{i}-1\right)\left(1-e^{-\frac{i}{n}}\right) \leq 0.$$

Hence $\left(\frac{n}{x}-1\right)\left(1-e^{-\frac{x}{n}}\right)$ decreases for $x>0$.

It follows that for $x \leq \frac{n}{2}$ we have

$$\mathbb{E}\left[I^{(t)} \mid I^{(t-1)}=x\right] \geq x\left(1+\frac{1}{\sqrt{e}}\right)$$

By the law of total expectation then

$\mathbb{E}\left[I^{(t)}\right] \geq \left(1+\frac{1}{\sqrt{e}}\right)^t$ as long as $\left(1+\frac{1}{\sqrt{e}}\right)^t \leq \frac{n}{2}$, that is

$$t=\mathcal{O}(\log n).$$

Back to PUSH

Law of Total Expectation

$$\mathbb{E}\left[\mathbb{E}\left[X \mid Y\right]\right] = \mathbb{E}\left[\sum_x x \Pr\left(X = x \mid Y\right)\right]$$

$$= \sum_y \left(\sum_x x \Pr\left(X = x \mid Y = y\right)\right) \Pr\left(Y = y\right)$$

$$= \sum_y \sum_x x \Pr\left(X = x,\ Y = y\right)$$

$$= \sum_x x \sum_y \Pr\left(X = x,\ Y = y\right)$$

$$= \sum_x x \Pr\left(X = x\right) = \mathbb{E}\left[X\right]$$

Back to PUSH

Bac to PULL

Upper Bound on Uninformed Agents

The uninformed nodes $U^{(t)}$ are

$$\mathbb{E}\left[U^{(t)} \,|\, U^{(t-1)} = u\right] = u - u\left(1 - \left(1 - \frac{1}{n}\right)^{(n-u)}\right).$$

Again, with $1 - y \leq e^{-y}$ we have

$$\mathbb{E}\left[U^{(t)} \,|\, U^{(t-1)} = u\right] \leq u - u\left(1 - e^{\frac{u}{n}-1}\right)$$

and for $u \leq \frac{n}{2}$ we have

$$\mathbb{E}\left[U^{(t)} \,|\, U^{(t-1)} = u\right] \leq u\left(1 - \frac{1}{\sqrt{e}}\right).$$

Expectation Requirement for Concentration via CB

Let X_i be independent r.v.s with $|X_i| \leq M$, then

$$\Pr\left(\left|\sum_i X_i - \mathbb{E}\left[\sum_i X\right]\right| \geq \Delta\right) \leq 2e^{-\frac{\Delta^2}{2\left(\sum_i \mathbb{E}[X_i^2] + \frac{M}{3}\Delta\right)}}$$

Since we want "w.h.p." we need that on the r.h.s. $e^{\text{quantity}} \leq \frac{1}{n}$, that is

$$\text{quantity} = -\frac{\Delta^2}{2\mathbb{E}\left[\sum_i X_i\right] + \frac{2}{3}\Delta} \leq \log\frac{1}{n} = -\log n$$

It is not hard to see that this is true for $\mathbb{E}[\sum_i X_i] = \Omega(\log n)$

Calculations for =>This Slide<=

We have

$$\text{quantity} = -\frac{\Delta^2}{2\mathbb{E}\left[\sum_i X_i\right] + \frac{2}{3}\Delta} \leq \log\frac{1}{n} = -\log n$$

that is $3\Delta^2 - 2\log n\Delta - 6\log n\mathbb{E}\left[\sum_i X_i\right] \geq 0$, which requires that

$$\Delta \geq \frac{1}{3}\log n\left(1 + \sqrt{1 + 18\frac{\mathbb{E}\left[\sum_i X_i\right]}{\log n}}\right)$$

or more simply

$$\Delta \geq 2\sqrt{\mathbb{E}\left[\sum_i X_i\right]\log n}$$

But we also need $\mathbb{E}\left[\sum_i X_i\right] \geq \Delta$, so $\mathbb{E}\left[\sum_i X_i\right] \geq 4\log n$.