Natural Distributed Algorithms

- Lecture 0 -
Introduction

Emanuele Natale
CNRS - UCA

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

This course

Course webpage: nda.enatale.name

How you will be graded:
At the end of each lecture I will propose possible projects.
You can pick one of them or propose your own idea.

Informal prerequisites:
- Discrete probability
- Distributed Algorithms
- Linear Algebra

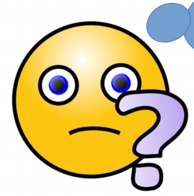

"What **if I miss some background**?"
Then many parts of lectures will be hard BUT many parts are still accessible, and you have to choose a final project with adequate necessary background.

Strong advice:
Do the project by the end of January
(much easier to get feedback)

What are (not) Natural Distributed Algorithms (NDA)

Natural Algorithms are NOT Natural Computing

Heuristics that take **inspiration from Nature** for the development of novel problem-solving techniques

Example:
Ant Colony Optimization Algorithms

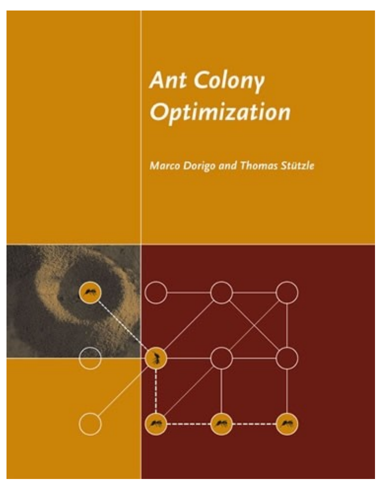

Instead:
Natural Algorithms are **algorithms that model biological** processes

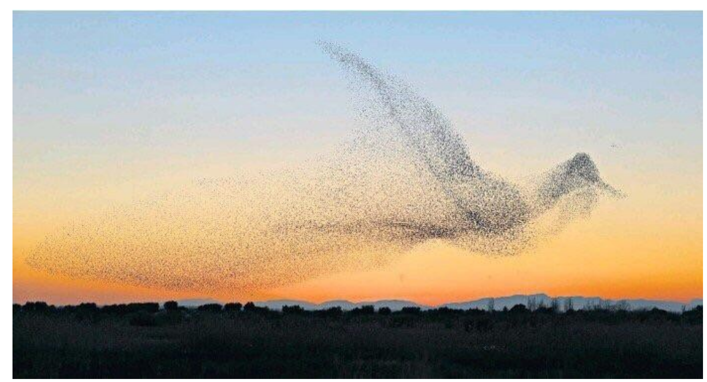

B. Chazelle, "Natural algorithms," in Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, 2009, pp. 422–431.

Collective Animal Behaviors as Complex Systems

A **computational lens** on how **global behavior** emerges from **simple stochastic interactions** among individuals

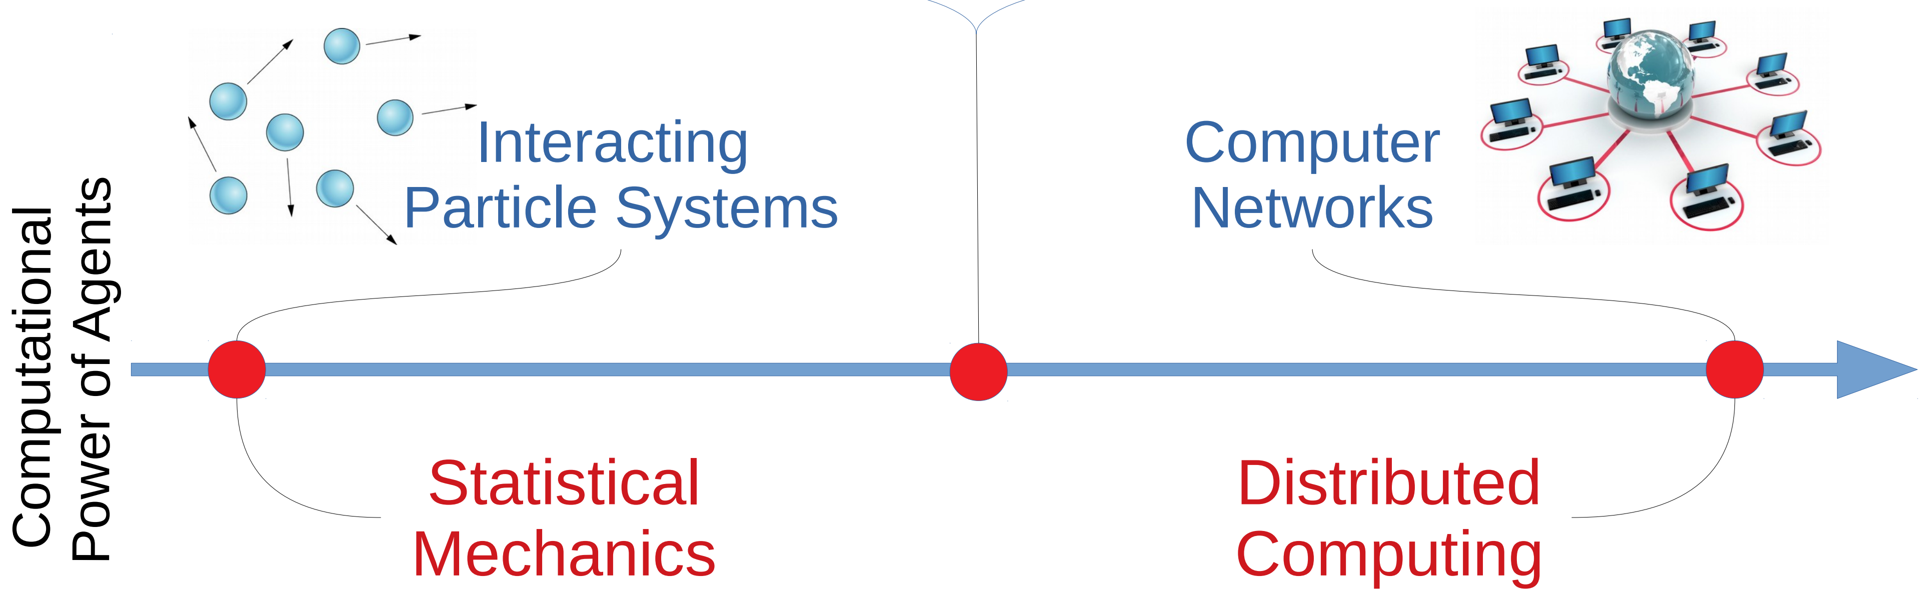

Natural Distributed Algorithms in Context

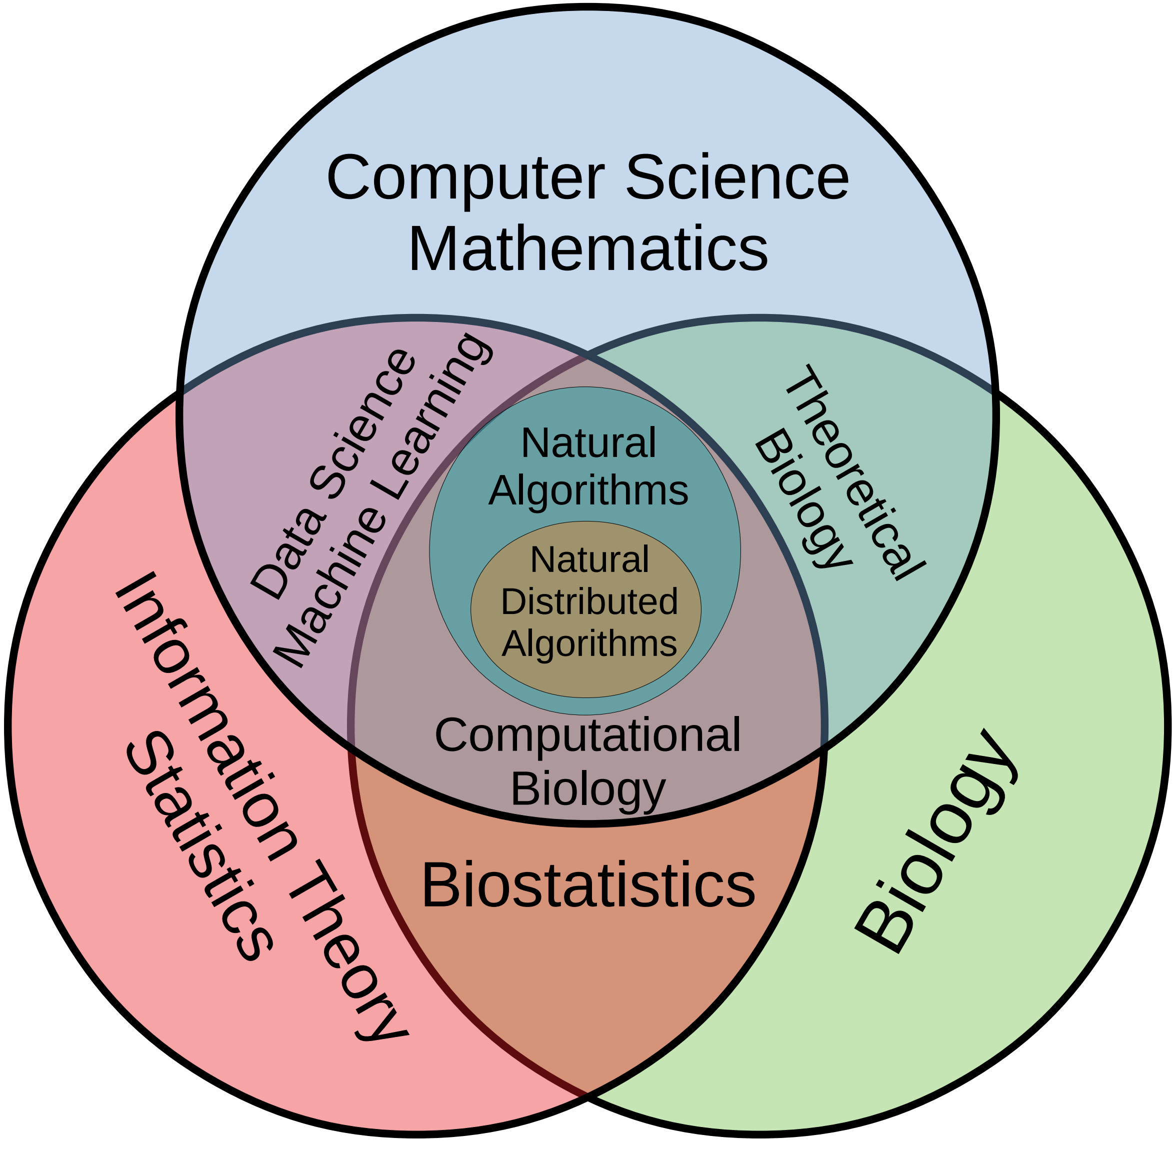

NDA: Why Now?

Biology:
New techniques for observing collective behaviors
(high-resolution cameras, fluorescence tagging, multi-electrode arrays...)

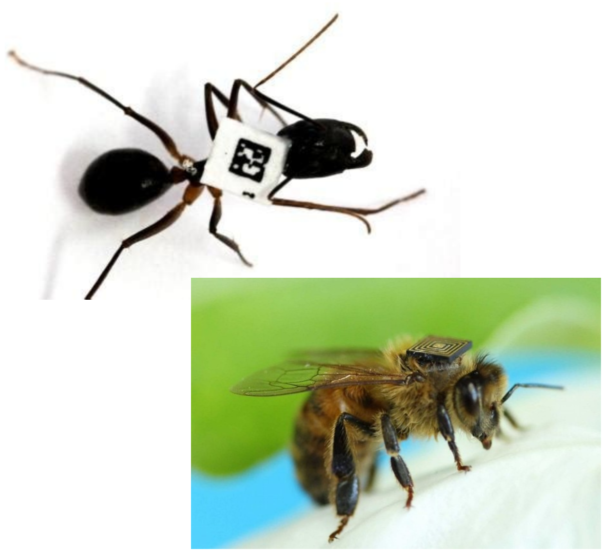

Distributed Computing:
New techniques for understanding weaker models
(dynamic networks, stochastic interactions, restricted memory and communication...)

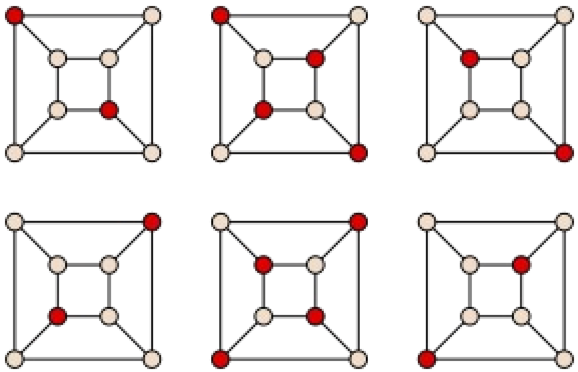

New CS Perspective to Biology

In biology the *model* specifies all aspect of the process at hand

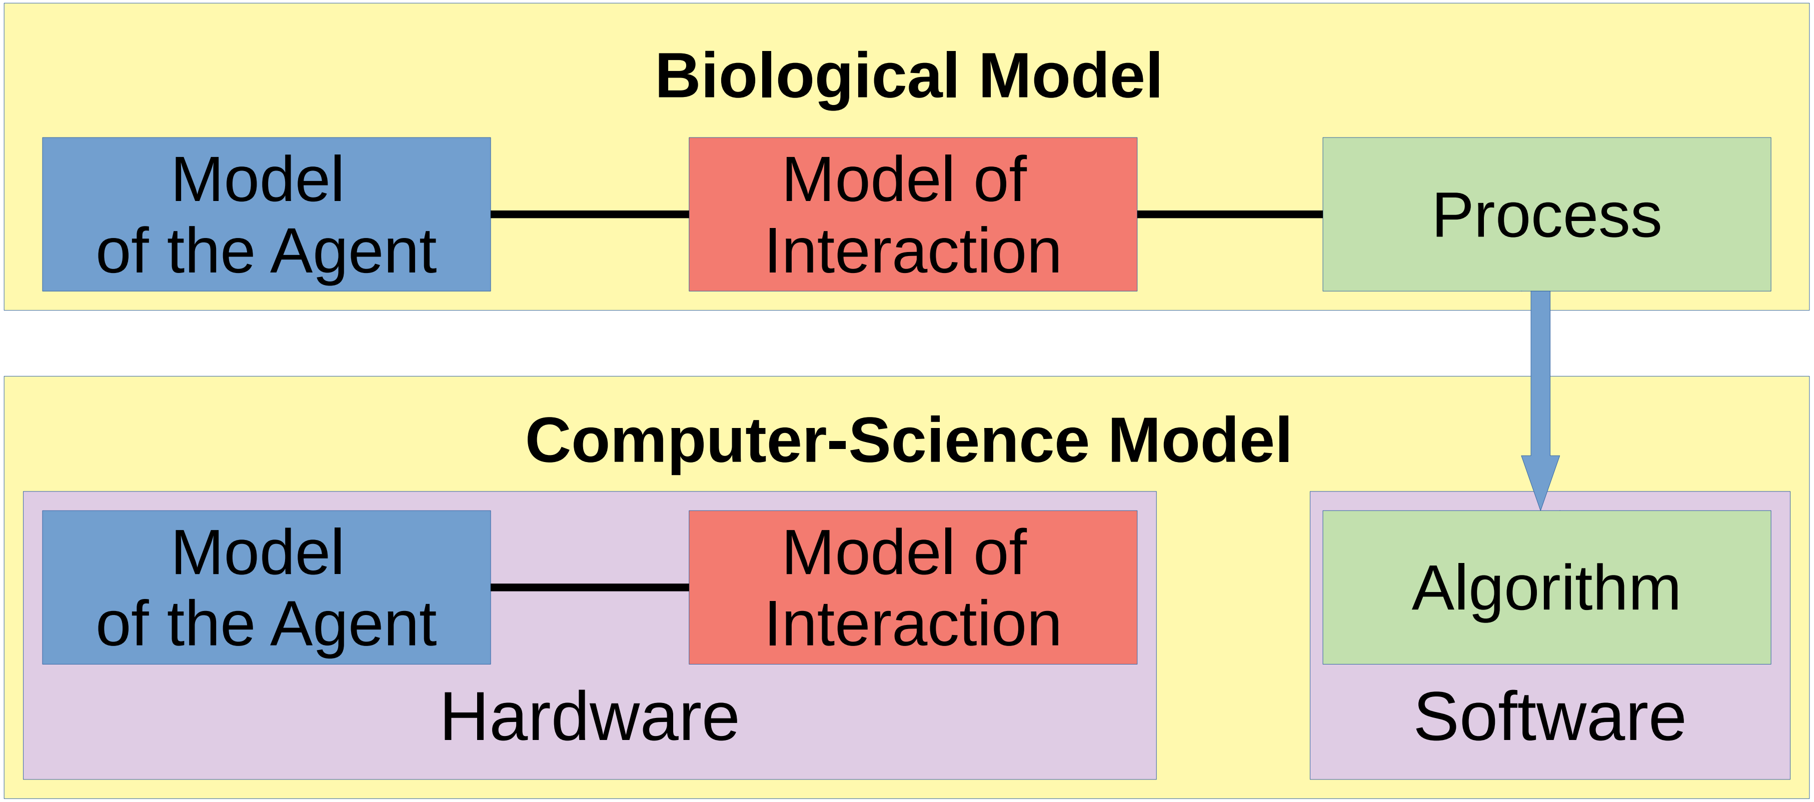

In CS the *model* only specifies constraints on the algorithm

Model vs Algorithm

	Known Model	Unknown Model
Known Algorithm	**Theoretical analysis of the algorithm:** • Chazelle, Bernard. 2009. “Natural Algorithms.” In Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, 422–431. Society for Industrial and Applied Mathematics. http://dl.acm.org/citation.cfm?id=1496817. • Bonifaci, Vincenzo. 2013. “Physarum Can Compute Shortest Paths: A Short Proof.” Information Processing Letters 113 (1–2): 4–7. https://doi.org/10.1016/j.ipl.2012.09.005.	**Finding a good abstraction of the model:** • (Example from Social Sciences) J. M. Kleinberg, “Navigation in a small world,” Nature, vol. 406, no. 6798, pp. 845–845, Aug. 2000.
Unknown Algorithm	**Computational complexity analysis** • Emek, Yuval, and Roger Wattenhofer. 2013. “Stone Age Distributed Computing.” In Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing, 137–146. PODC ’13. https://doi.org/10.1145/2484239.2484244. **Guessing the algorithm** • 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.	**Surmising** • Y. Afek, N. Alon, O. Barad, E. Hornstein, N. Barkai, and Z. Bar-Joseph, “A biological solution to a fundamental distributed computing problem,” Science, vol. 331, no. 6014, pp. 183–185, Jan. 2011. **Finding dependencies between parameters** • L. Boczkowski, E. Natale, O. Feinerman, and A. Korman, “Limits on reliable information flows through stochastic populations,” PLOS Computational Biology, vol. 14, no. 6, p. e1006195, Jun. 2018.

Feinerman, Ofer, and Amos Korman. 2013. “Theoretical Distributed Computing Meets Biology: A Review.” In Distributed Computing and Internet Technology, 1–18. LNCS 7753. Springer Berlin Heidelberg. http://link.springer.com/chapter/10.1007/978-3-642-36071-8_1.

Algorithm-Driven Experiment Design

Stage 1
Find abstract setting parametrized by x that can be experimentally tested

Stage 2
Analyze the model and obtain theoretical trade-offs between x and the algorithm efficiency

Stage 3
Measure experimentally the efficiency of the biological system

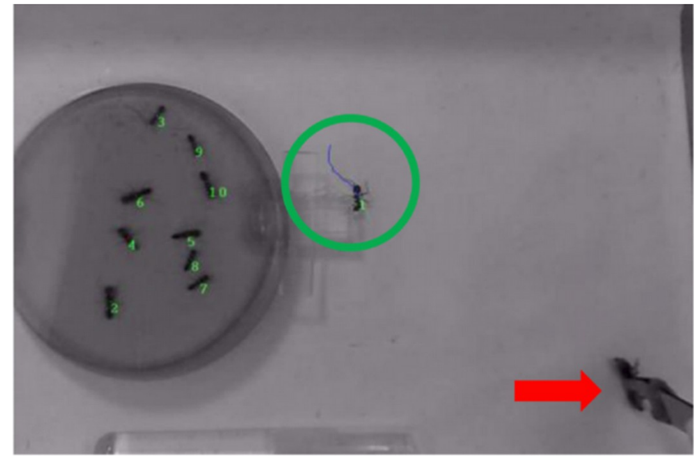

Theorem.
Rumor spreading takes $\tilde{\Theta}(n)$

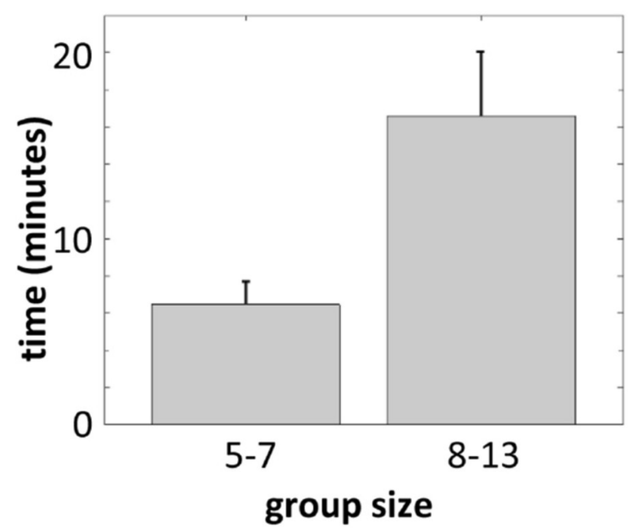

Project Idea

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

- **Feinerman, Ofer, and Amos Korman. 2013. “Theoretical Distributed Computing Meets Biology: A Review.” In Distributed Computing and Internet Technology, edited by Chittaranjan Hota and Pradip K. Srimani, 1–18. Lecture Notes in Computer Science 7753. Springer Berlin Heidelberg. http://link.springer.com/chapter/10.1007/978-3-642-36071-8_1.**
- Karp, Richard M. 2011. “Understanding Science Through the Computational Lens.” Journal of Computer Science and Technology 26 (4): 569–77. https://doi.org/10.1007/s11390-011-1157-0.
- Navlakha, Saket, and Ziv Bar-Joseph. 2011. “Algorithms in Nature: The Convergence of Systems Biology and Computational Thinking.” Molecular Systems Biology 7 (November): 546. https://doi.org/10.1038/msb.2011.78.
- **———. 2014. “Distributed Information Processing in Biological and Computational Systems.” Communications of the ACM 58 (1): 94–102. https://doi.org/10.1145/2678280.**
- The website http://algorithmsinnature.org/
- Works appeared in the Biological Distributed Algorithms Workshop.

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

Trailer for Guest Lecture on November 18th

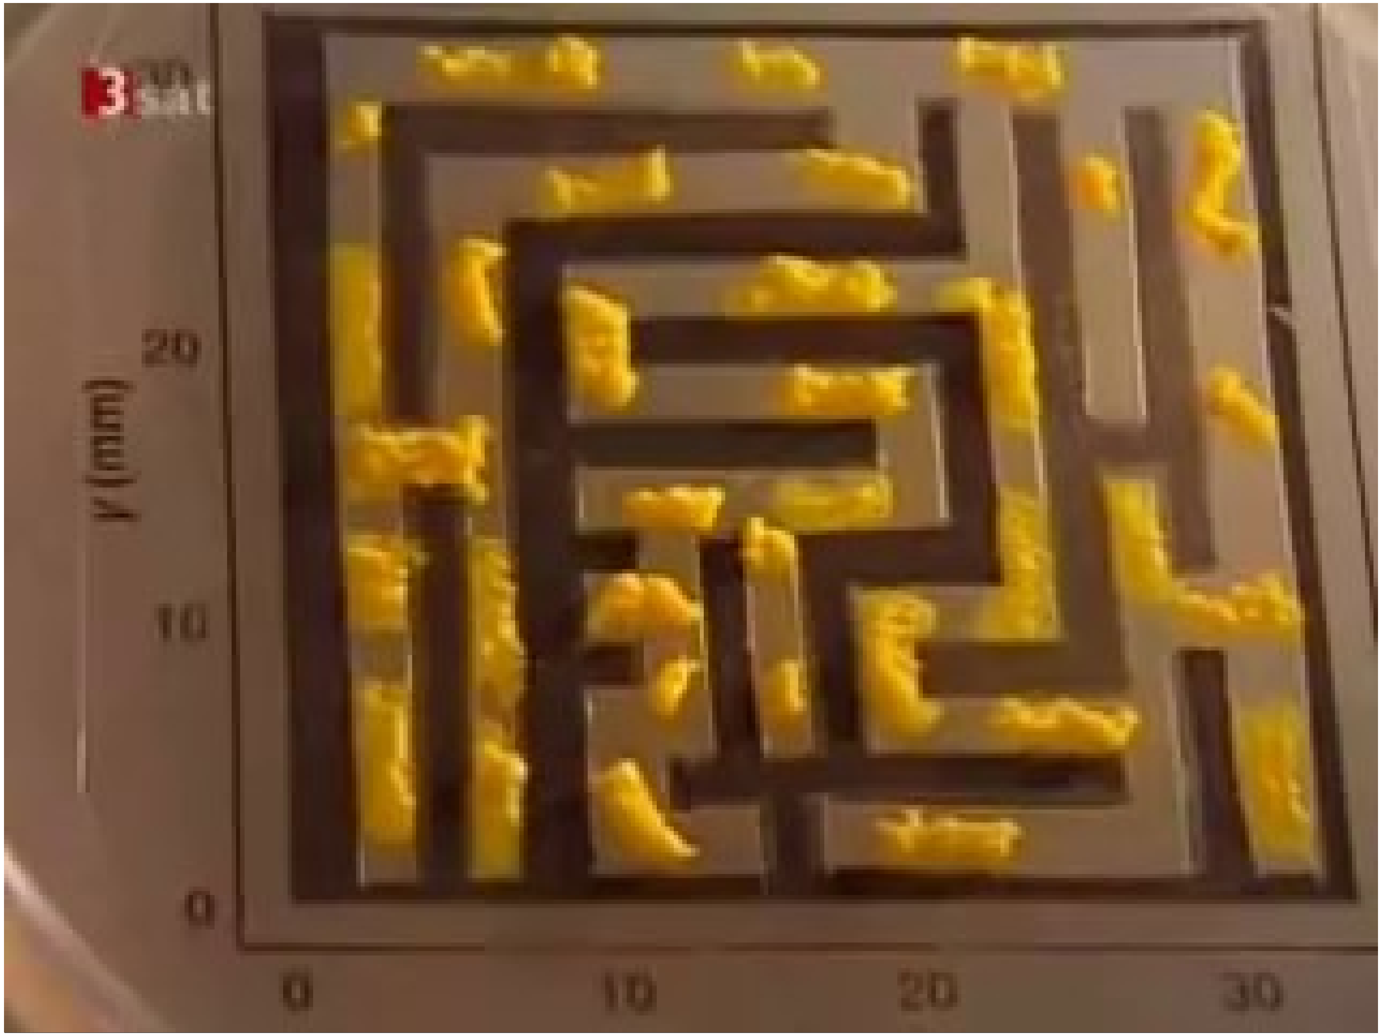

Dott. Vincenzo Bonifaci, co-author of several of the main algorithmic results on the **Physarum Dynamics**, will introduce the topic and present a new model.