

UiO : Department of Technology Systems
University of Oslo

TEK5010/9010 - Multiagent systems 2022
Lecture 1

Introduction

Jonas Moen

How to follow the course?

1. Physical lectures/exercises at ITS

2. Live streaming on Zoom (meeting links posted on the course web page)

3. Recorded videos available after each lecture (also posted on the web page but only available for registered students of the course)

My background

- Scientist FFI, leader swarm group
  Research in MultiX (Multi-agent systems, Multi-function systems, Multi-objective optimization)
- Associate professor UiO/ITS, teacher and supervisor
- Formal academic background
  - PhD in evolutionary computation, UiO/Robin 2013
  - Master in economics, NHH Bergen 2002
  - Master in physics, UiB 1999

What are Multi-Agent Systems (MAS)?

“Multiagent systems are systems composed of multiple interacting computing elements, known as agents.” [Wooldridge, 2009]

Agents are capable of *autonomous* actions satisfying design objectives:

1. *Reactive (or simple) agents* can produce *emergent properties* usually modelled by *Swarm Intelligence* (SI).
2. *Interactive (or intelligent) agents* engage in *social activities* like cooperation, coordination, negotiation, etc usually described using *Game Theory* (GT).

Reserach goal of multiagent systems

To connect micro scale behaviour with macro scale (often emergent) properties/effects, and vice versa.

Outline of the course

Part I - Agents, communication and cooperation

Part II - Swarm intelligence

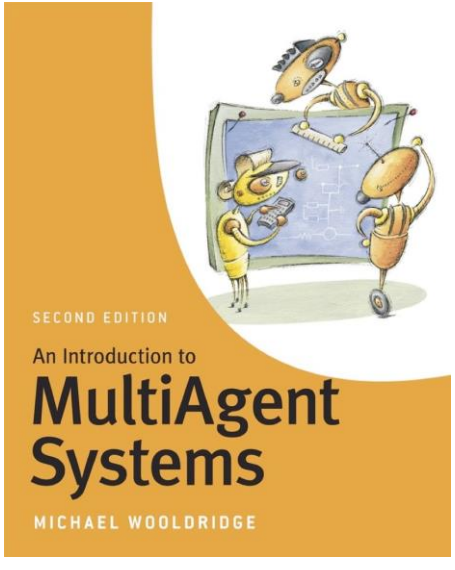
Part III - Game theory

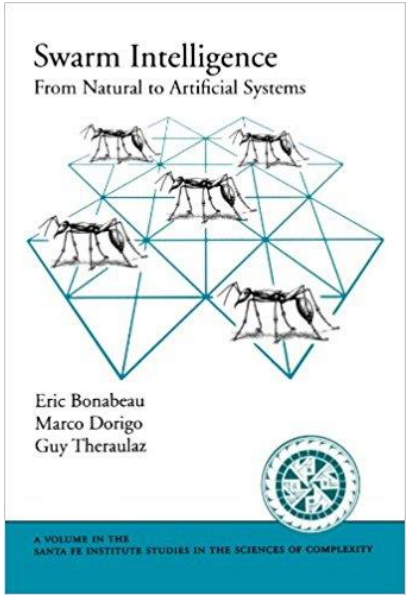
Weekly lecture, 2 hours theory and 1 hour exercises.

To pass course, 2 ‘obliger’ + exam.

Books in the course

«Multiagent Systems» (MAS)
Michael Wooldrigde, 2009

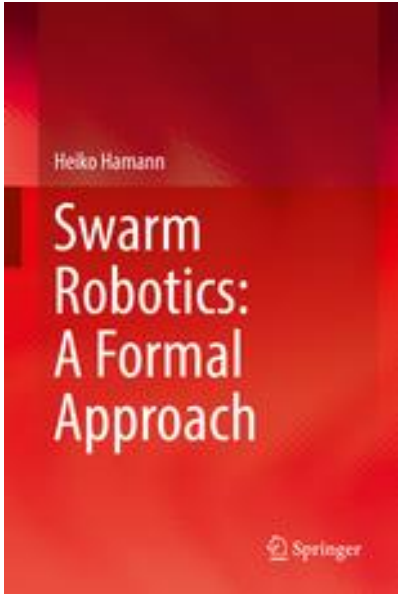




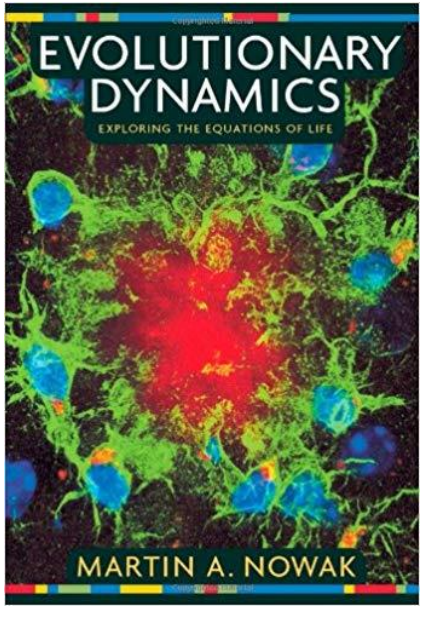
«Swarm Intelligence» (SI)
Dorigo et *al*., 1999

Books in the course

«Swarm Robotics» (SR)
Heiko Hamann, 2018



TEK9010 Extra PhD curriculum



«Evolutionary Dynamics» (ED)
Martin Nowak, 2006

Highlights lecture 1 – Introduction*

- What is a multiagent system?
  - Key concepts
  - In relation to trends in computing
  - When to use and apply
- A brief history of multiagent systems
- Some examples of multiagent system research and application

*Wooldridge, 2009: chapter 1, 9, 10, preface and appendix A.

Trends in computing

5 important and continuing trends [Wooldridge, 2009]

1. Ubiquity
2. Interconnection
3. Intelligence
4. Delegation
5. Human-orientation

Trends in computing

1. Ubiquity

   The reduced cost of computation (i.e. Moore’s Law) has made it possible to introduce processing power in almost all devices and places.

Trends in computing

2. Interconnection

   Computers are networked into large distributed systems (i.e. the Internet)

Trends in computing

3. Intelligence

   The increased complexity of tasks that we are able to automate and computerize in terms of research into algorithm design (i.e. AI)

Trends in computing

4. Delegation

   Implies that we give control to computer systems (e.g. safety-critical tasks in aerospace and energy)

Trends in computing

5. Human-orientation

   Human-machine programming and interfacing; from assembler languages, through procedural abstraction, via abstract data types, to object-oriented programming and, now most recently, agent-based systems.

Major challenges in software development

1. With respect to ubiquity, interconnection and AI, how do we develop techniques that exploit this massive processing power without explicitly telling the computer what to do? (e.g. ~10^{10} processing units)

   Can we make an optimal system that is not hard-coded? Do we need an alternative software paradigm in order to tap into this immense computing power?

Major challenges in software development

2. With respect to delegation and human-machine interaction, how do we build computer systems that can act effectively on our behalf? They must be
   - Independent and autonomous
   - Represent our best interests

In a dynamical world of other agents and humans (on the net, in cyber-physical systems, robots, autonomous vehicles, and in the social sciences and biological domain).

Multiagent systems

A relatively new field in computer science trying to answer these 5 trends in software development

1. An agent is a computer system capable of independent action on behalf of its user or owner.
2. A multiagent system consist of many such agents interacting with each other (through some network or sensor system)

Agents need skills and abilities to cooperate, coordinate and negotiate with each other on behalf of their users.

Key challenges with multiagent systems

1. The agent design problem (the micro level)
   How do we build agents that are capable of independent, autonomous action in order to successfully carry out tasks that we delegate to them?
2. The social design problem (the macro level)
   How do we build agents that are capable of interacting with other agents in order to successfully carry out the tasks that we delegate to them, especially when the agents do not share common goals or intentions?

Key challenges with multiagent systems

A truly successful multiagent system makes an explicit connection between the autonomous micro level agents and the macro level modelling of the complex system.

Social abilities of multiagent systems

1. How can cooperation emerge in societies of self-interested agents?
2. How can self-interested agents recognize when their desires and beliefs (goals and actions) conflict and avoid resorting to conflict?
3. How can autonomous agents coordinate their activities so as to cooperatively achieve goals?

Social abilities of multiagent systems

4. What sort of common language can agents use to communicate their desires and beliefs?
5. How can we support this kind of cooperative system when agents have different software and hardware platforms?

And most importantly, multiagent systems are based on artificial computational entities, that often are conceptually inspired by social human systems and biologic systems.

A brief history of multiagent systems research

A history of agents

The notion of ‘agent’ is evident in early AI literature,
e.g. the Turing test and the term AI coined by John McCarthy at the famous Dartmouth workshop in 1956.

A brief history of multiagent systems research

1940-1980: Symbolic AI

Until the mid 1980s the AI community was mainly occupied with formal logic and deductive reasoning systems known as symbolic AI. STRIPS [Fikes and Nilsson, 1971] is prominent example of this era.

AI researchers were working in subfields like learning, planning, reasoning and so forth, but not really interested in a ‘holistic’ approach to agency.

A brief history of multiagent systems research

1980s: ‘new AI’

In the mid 80s the ‘new AI’, also called ‘behavioural AI’ and ‘reactive AI’ were introduced in opposition to symbolic AI.

1. Inspiration from biology, spawning new research fields like evolutionary computation, Alife, swam intelligence, etc.
2. Realistic scenarios were adopted in research and development.
3. Hybrid agents combining deliberate reasoning and reactive behaviour in same architecture.

A brief history of multiagent systems research

A history of multiagent systems

Multiagent systems, as a research field of its own, did not gain widespread popularity and attention in the AI community until early 1980s.

A brief history of multiagent systems research

1960-1970: The blackboard model [Newell, 1962]

1970s: The actor model of computation, [Hewitt, 1977]

- Actors are reactive but not proactive
- Closely connected to the development of object-oriented languages
- The Contract Net [Smith, 1977] introduced the economic metaphore and negotiations in multiagent systems

A brief history of multiagent systems research

1980-1990: The ‘distributed AI’ paradigm

- Spawned several workshops, conferences and publications
- The concept of self-interested agents introduced game theory to multiagent systems, [Rosenschein and Genesereth, 1985].
- MACE (MultiAgent Computing Environment) pioneered the acquaintance model, [Gasser et *al*., 1987]
- ARCHON (Architecture for cooperating heterogenous online systems) as first real industrial application of agent systems. Financed by EU in late 1980s, 14 partners in 9 countries.

A brief history of multiagent systems research

1990->: The Internet

- In 1990s interest in multiagent systems grew very rapidly.
- E-commerce represents a natural and lucrative application domain with the introduction of online auctions, shopping agents, personal web robots and now IoT.
- Standardization became a major issue and by 2000 FIPA emerged as the central standard of multiagent systems.
- Multiagent systems are now mainstream computer science.

A brief history of multiagent systems research

2010->: Autonomous mobile robots

In the beginning of the 2010s autonomous robots started to move out of the academic university laboratories and into real world field trials. This was due to reduced cost and increased availability of appropriate hardware (and SatNav).

A brief history of multiagent systems research

2010->: Autonomous mobile robots

- Autonomus systems take inspiration from the like of multiagent system research, evolutionary computation methods, swarm intelligence, etc.
- Produced massive interest in autonomous cars and in novel applications like swarm systems of drones and other vehicles, applied to transport, energy, agriculture, etc.
- Note that mobile agents are also software programs that can transmit themselves over the net and recommence execution at remote site.

Multiagent system research and applications

Some examples:

1. A software perspective on IoT
2. Modelling social behaviour: Hotelling's beach
3. What is cooperation? Nowak's spatial game
4. Modern swam robotics: Locust and Kilobots

A software perspective on IoT



Image: IoT Tech Expo

A software perspective on IoT

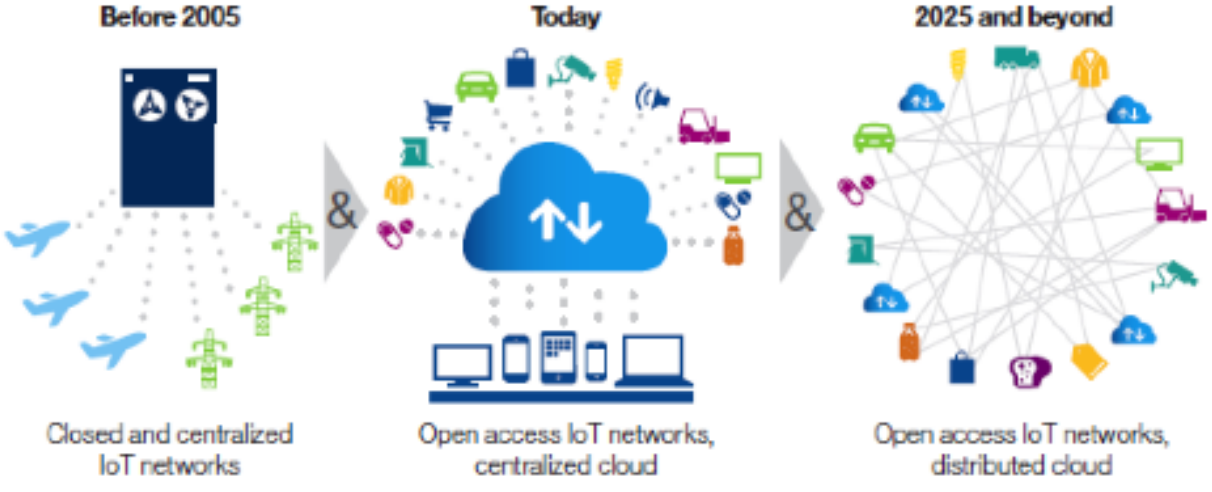


Image: IBM Big Data & Analytics Hub

A software perspective on IoT

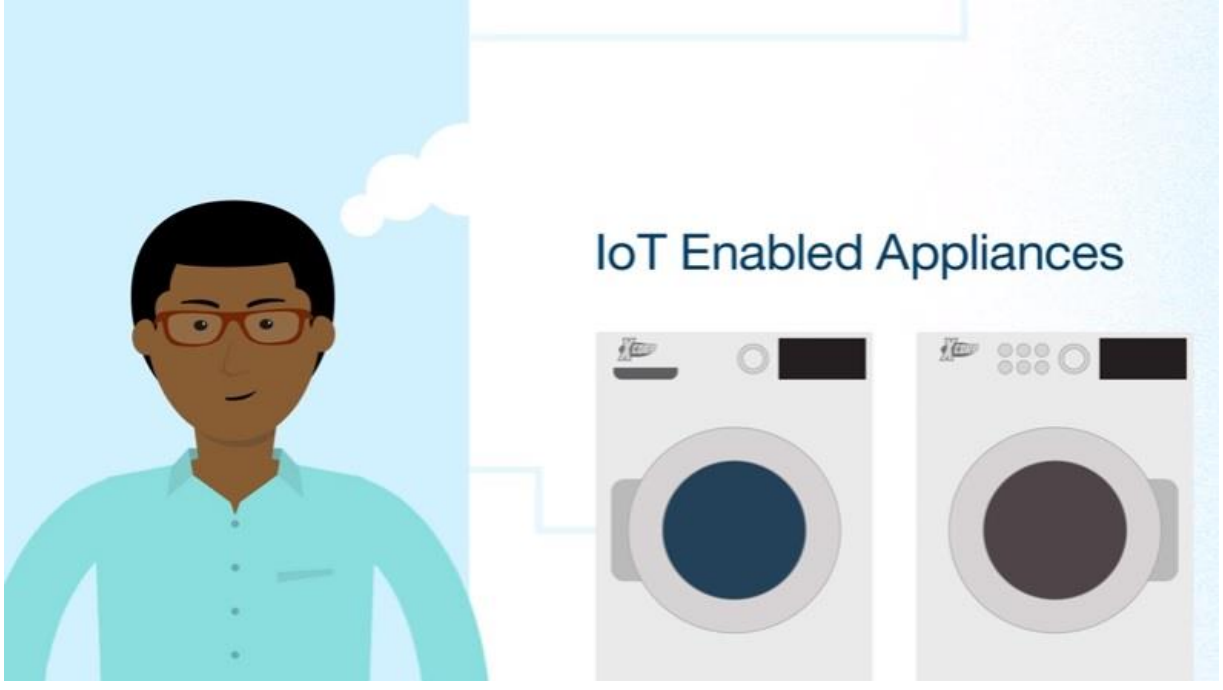


Image: IBM Big Data & Analytics Hub

A software perspective on IoT

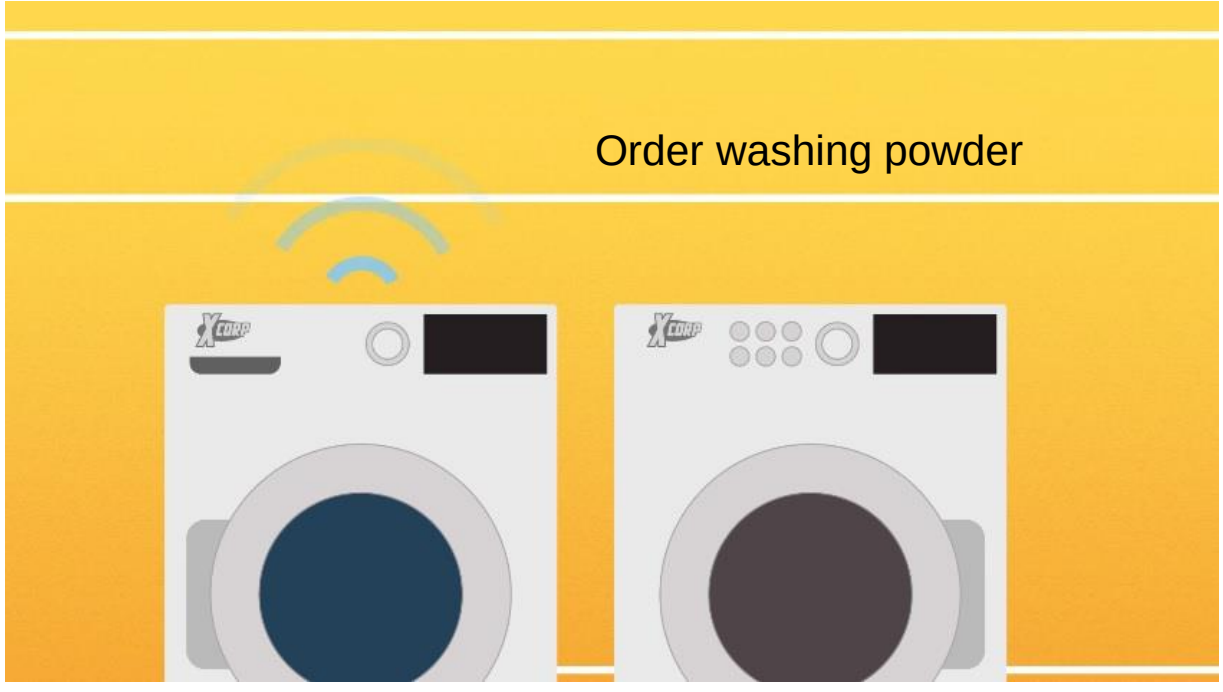
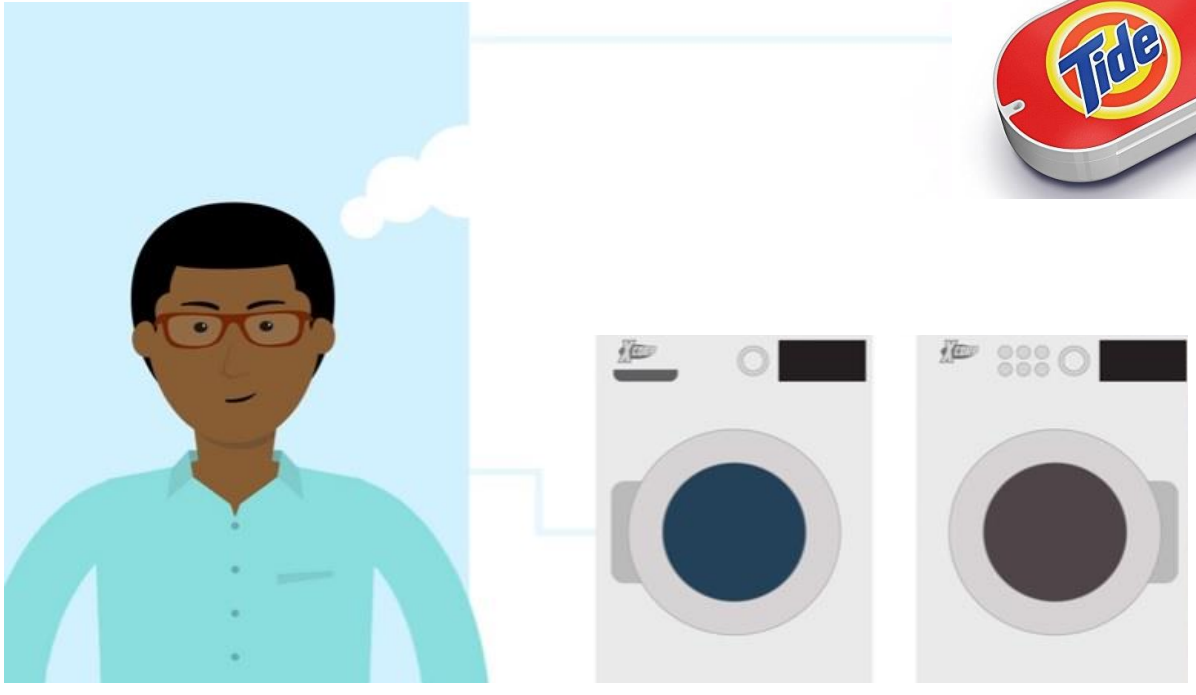


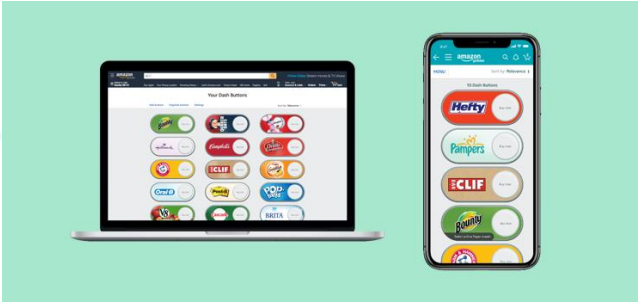
Image: IBM Big Data & Analytics Hub

A software perspective on IoT

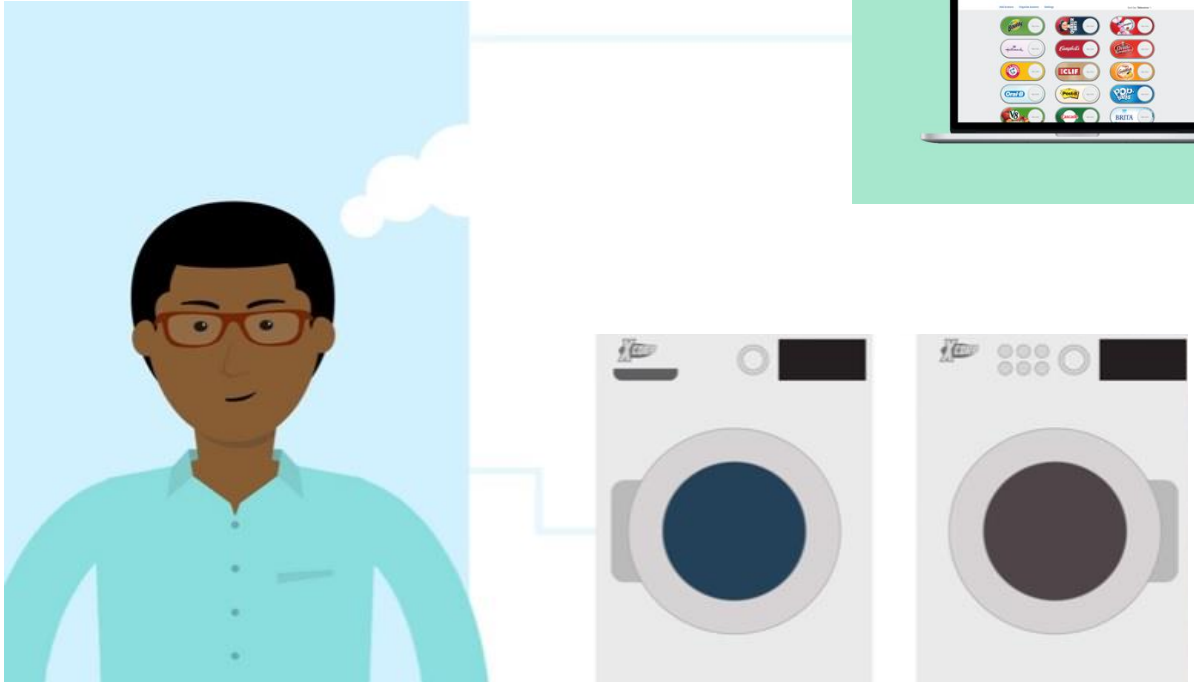


Images: IBM Big Data & Analytics Hub
Amazon Tide Push Button

A software perspective on IoT



2020



Images: IBM Big Data & Analytics Hub
Amazon Tide Push Button

A software perspective on IoT

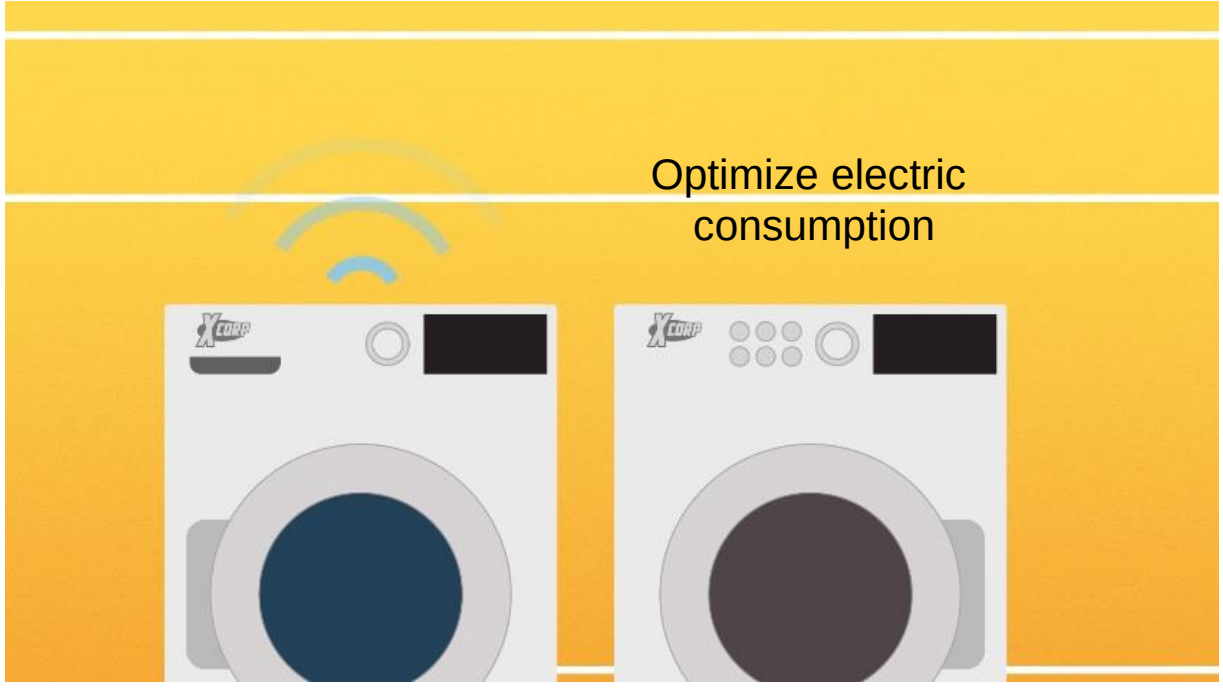


Image: IBM Big Data & Analytics Hub

A software perspective on IoT





Images: Aidon Power Grid Management
SmartGrid.gov, DoE

A software perspective on IoT





Images: Aidon Power Grid Managament
SmartGrid.gov, DoE

Modelling social behaviour: Clustering of services

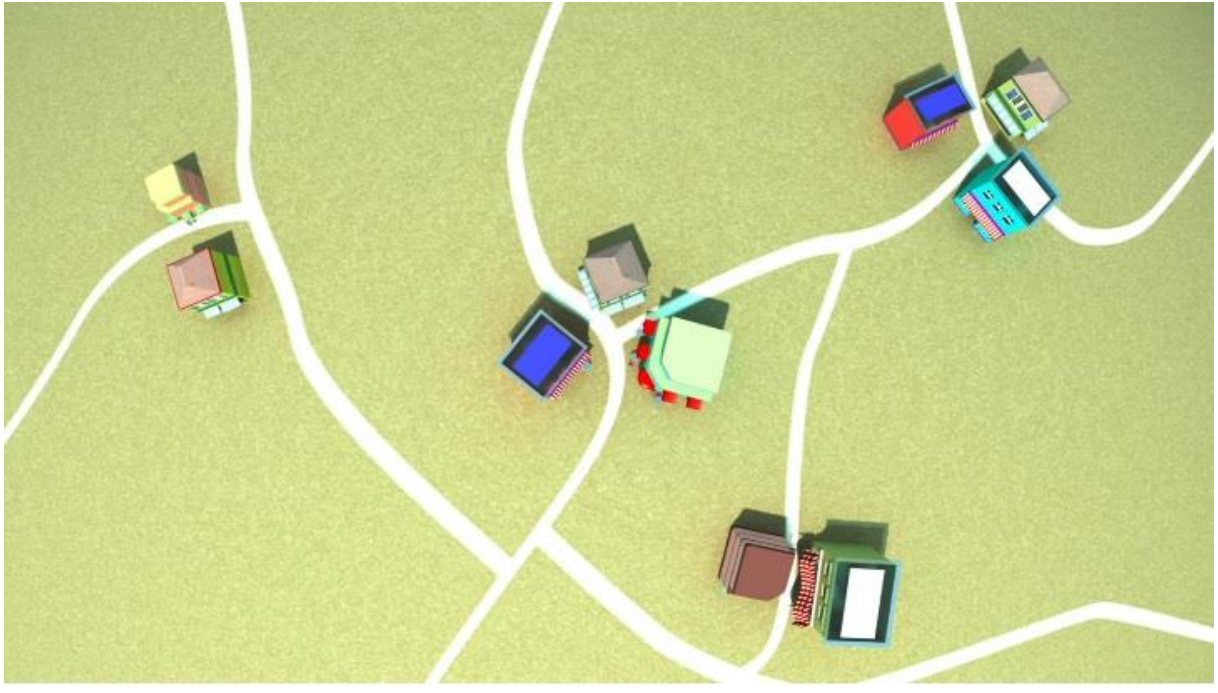


Image: TedEd

Modelling social behaviour: Hotelling's beach*



*Hotelling, «Stability In Competition», TEJ, 1929.

Image: TedEd

Modelling social behaviour: Hotelling's beach

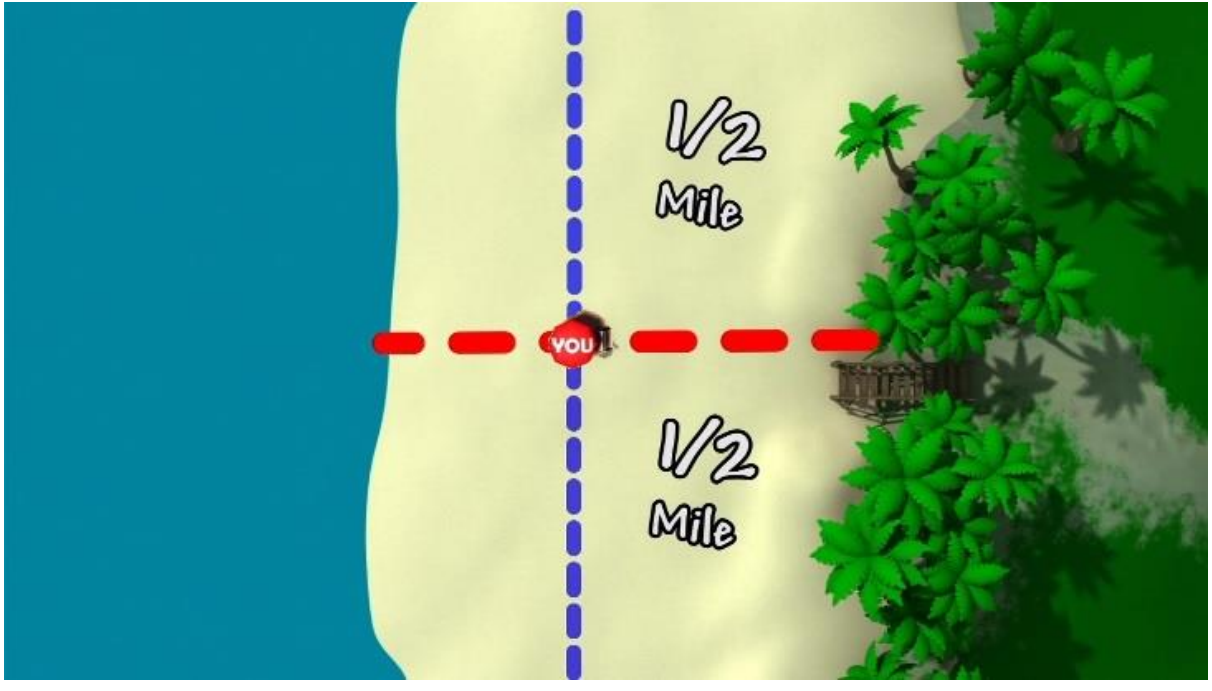


Image: TedEd

Modelling social behaviour: Hotelling's beach

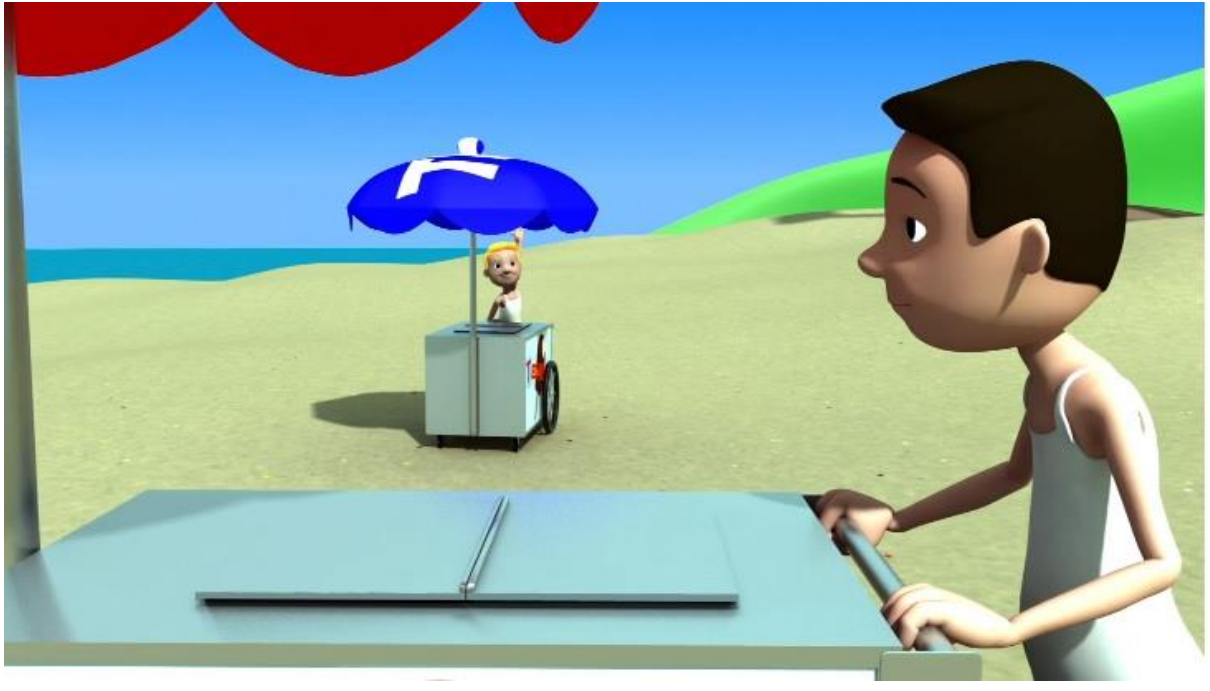


Image: TedEd

Modelling social behaviour: Hotelling's beach

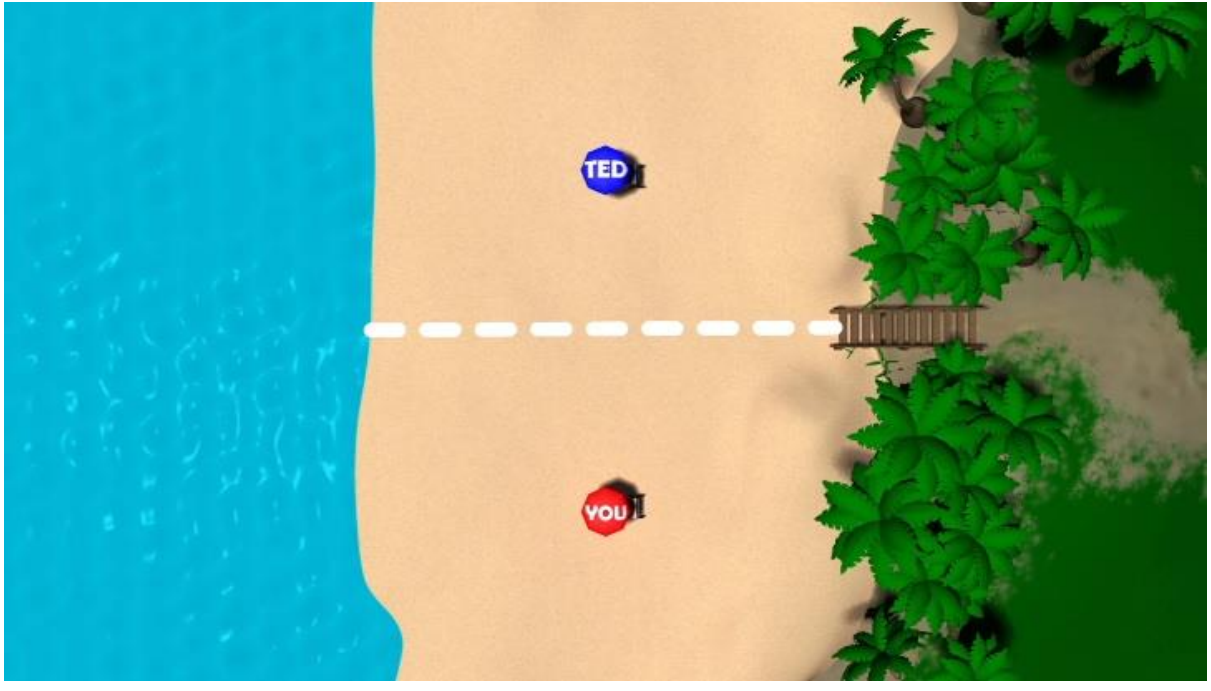


Image: TedEd

Modelling social behaviour: Hotelling's beach

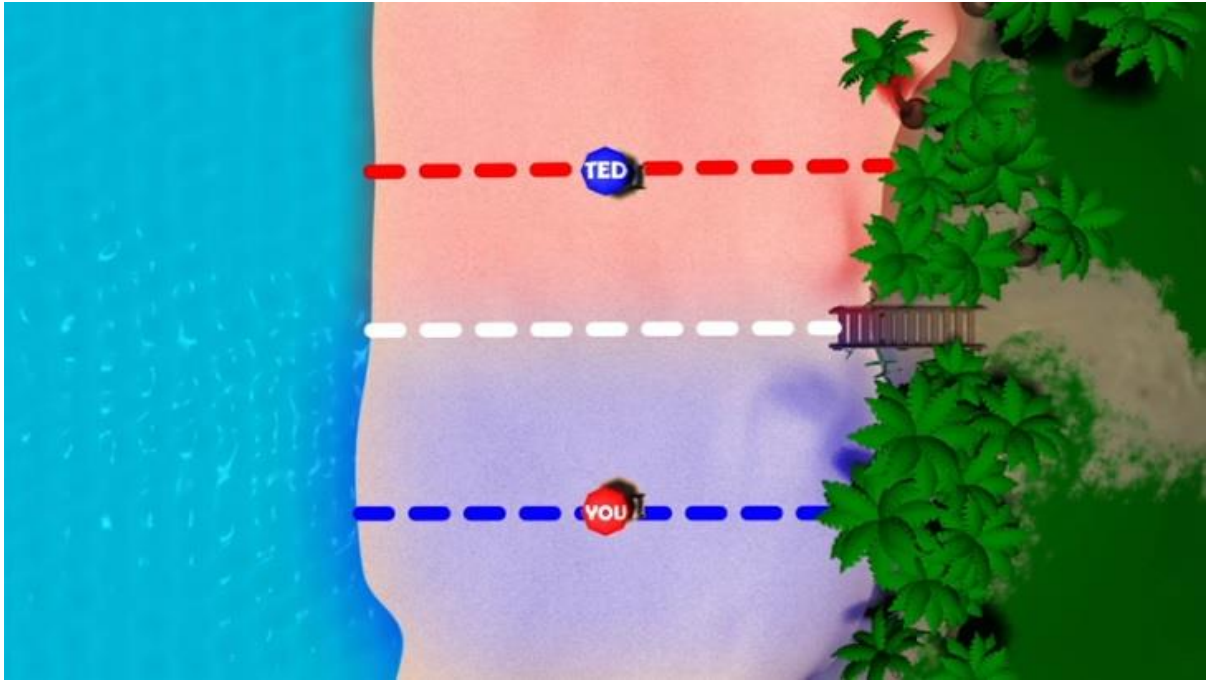


Image: TedEd

Modelling social behaviour: Hotelling's beach

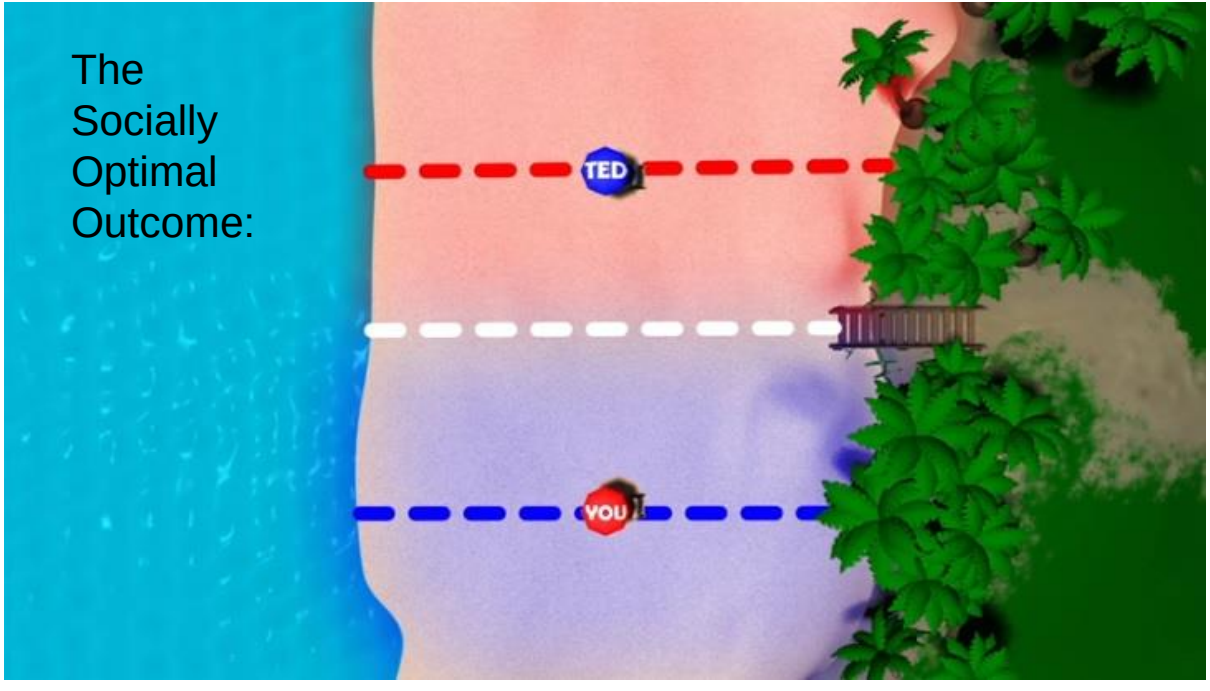


Image: TedEd

Modelling social behaviour: Hotelling's beach

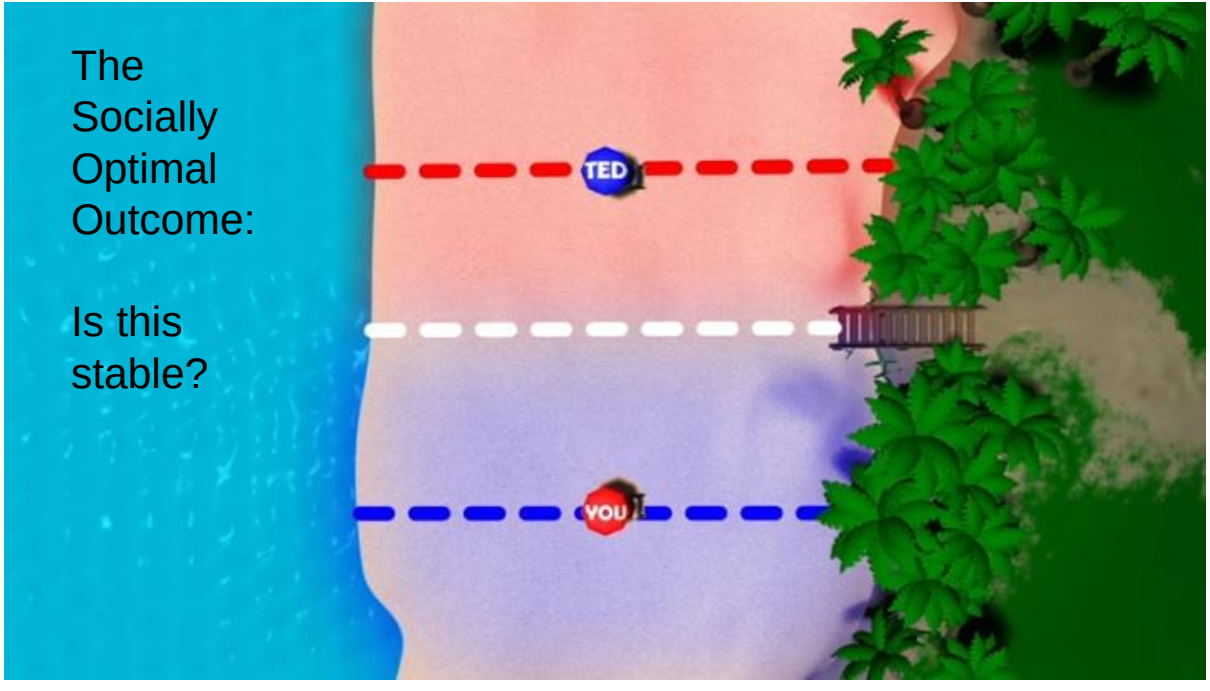


Image: TedEd

Modelling social behaviour: Hotelling's beach

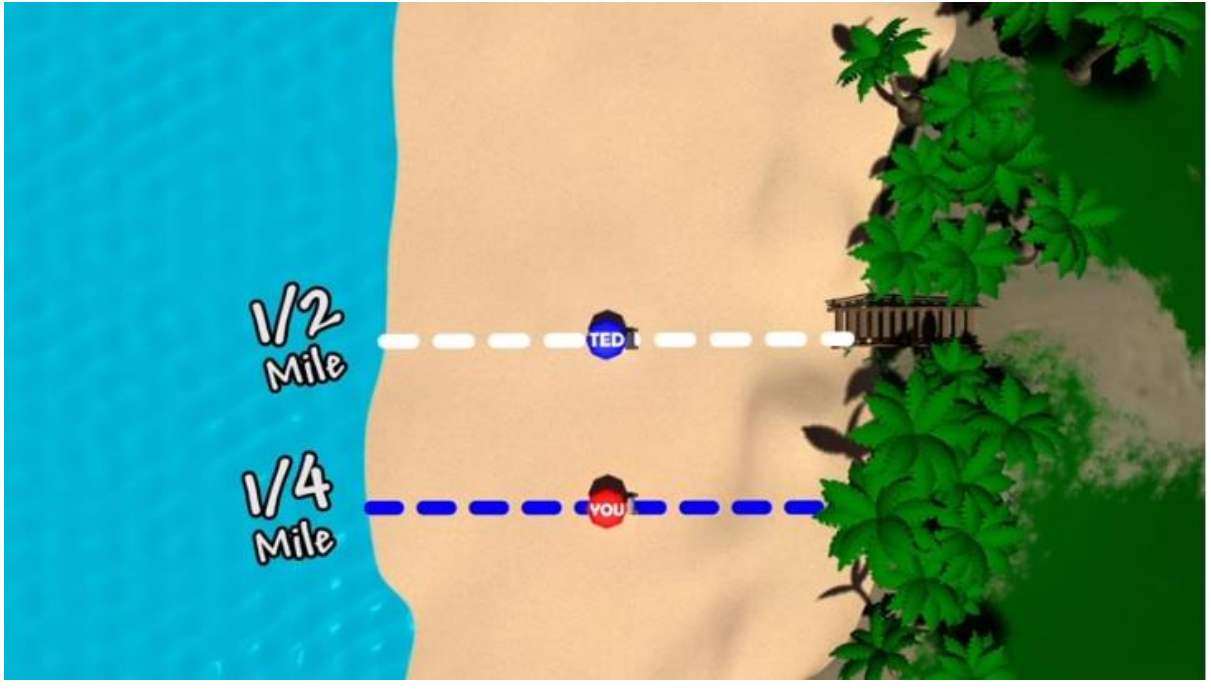


Image: TedEd

Modelling social behaviour: Hotelling's beach

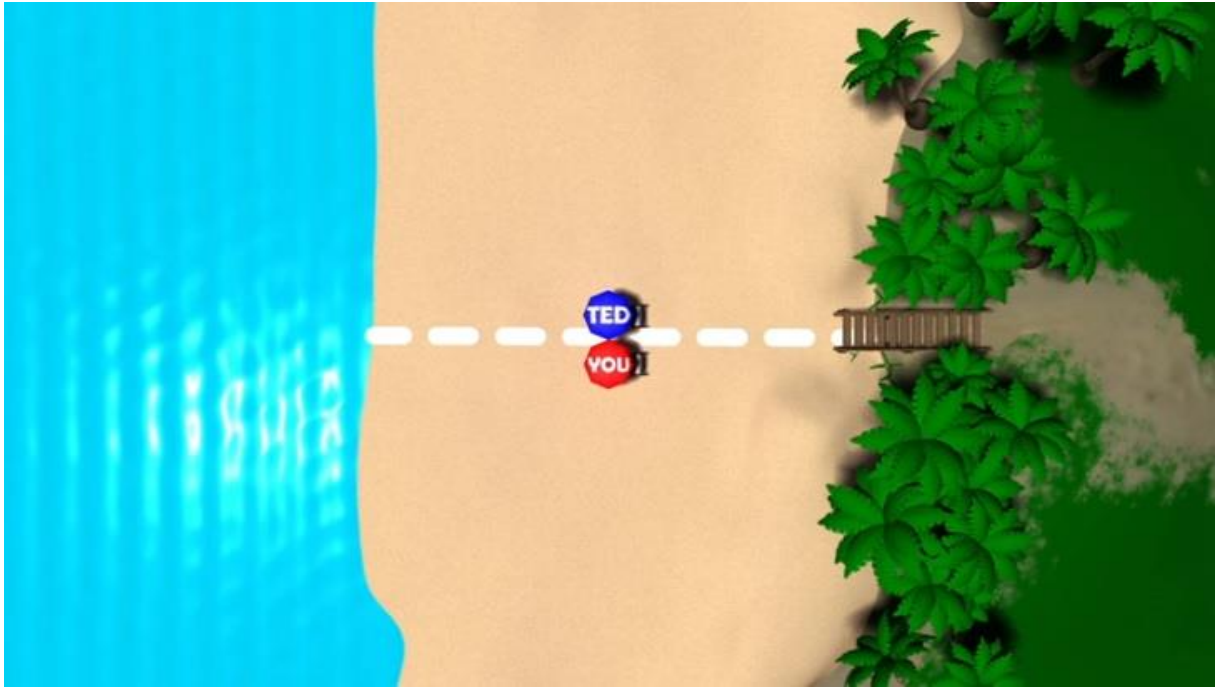


Image: TedEd

Modelling social behaviour: Hotelling's beach

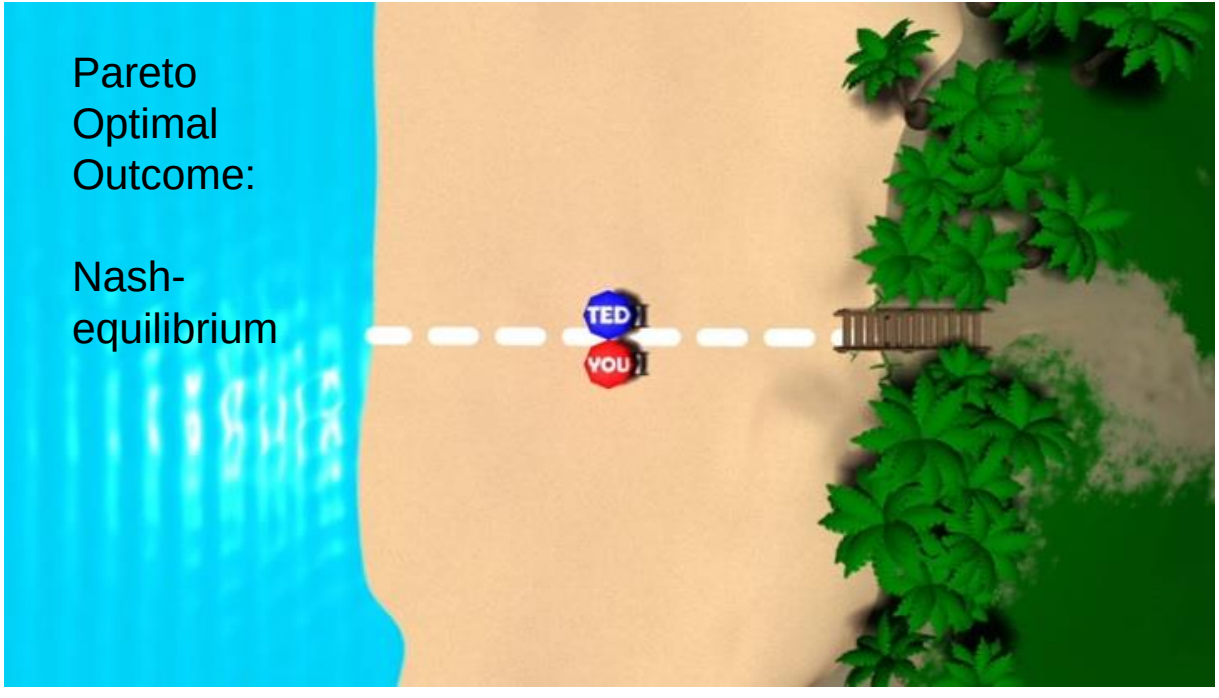


Image: TedEd

Modelling social behaviour: Hotelling's beach

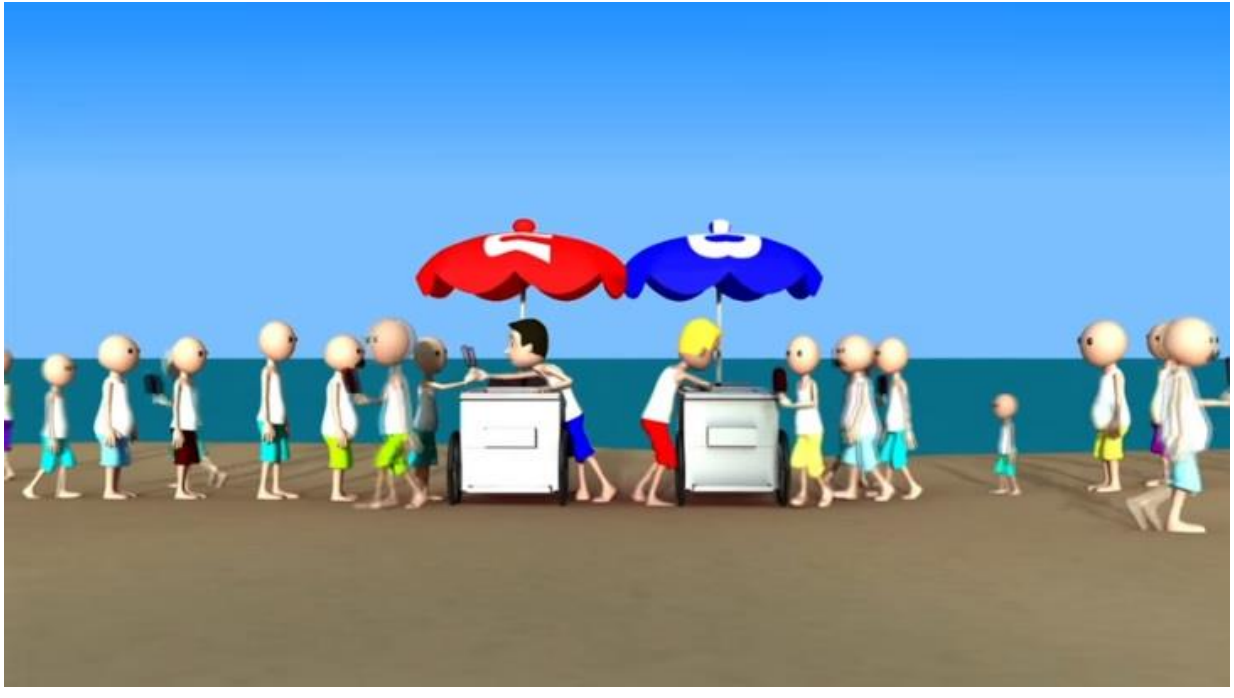
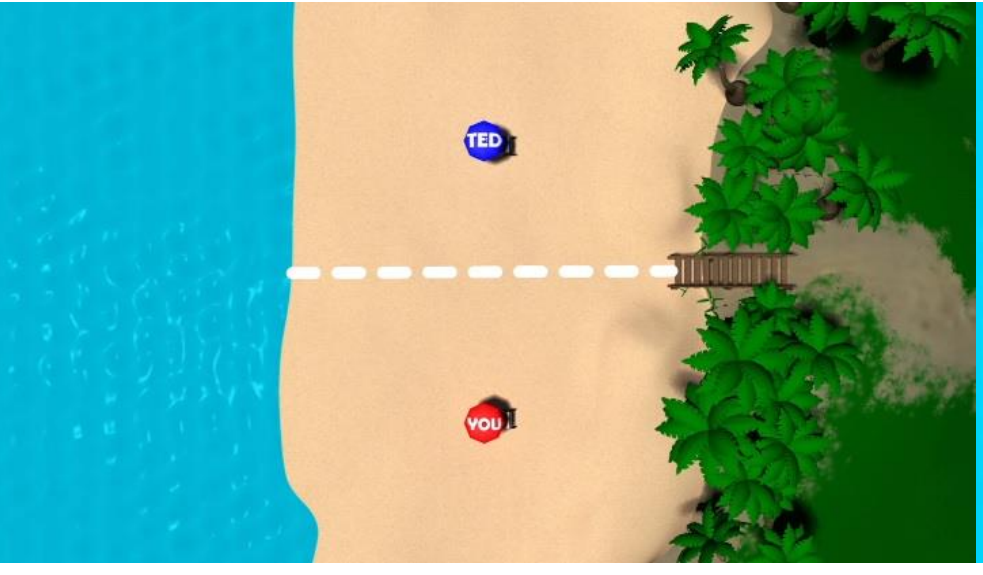


Image: TedEd

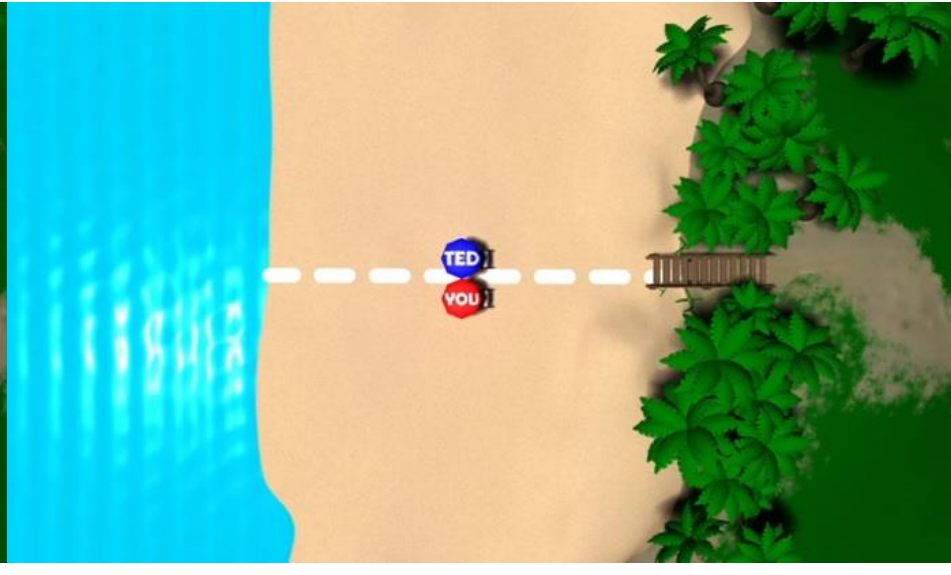
Modelling social behaviour: Clustering of services

Socially optimal solution:



Sales:1/2 Customer: 1/4

Pareto optimal solution:



Sales:1/2 Customer: 1/2

Modelling social behaviour: Clustering of services

Now, based on Hotelling's analysis, how would a set of mobile autonomous 4G base stations allocate in order to cover an ad-hoc crowd of people addicted to SoMe?

1. Self-interested robots would cluster in Nash-equilibrium
2. Altruistic robots would spread out to the social optimum

The really big question: is it possible to produce a socially optimal solution using rational self-interested agents?

What is cooperation? Nowak's spatial game*

The cooperation puzzle:

- How can cooperation be maintained when there is a clear advantage in defecting?
- How can cooperation be established when there is a clear advantage in defecting?

Some researchers, among them Nowak, believe that life itself presuppose cooperation (which is always under the threat of opportunism).

*Nowak et *al.*, "Evolutionary games and spatial chaos". Nature 1992

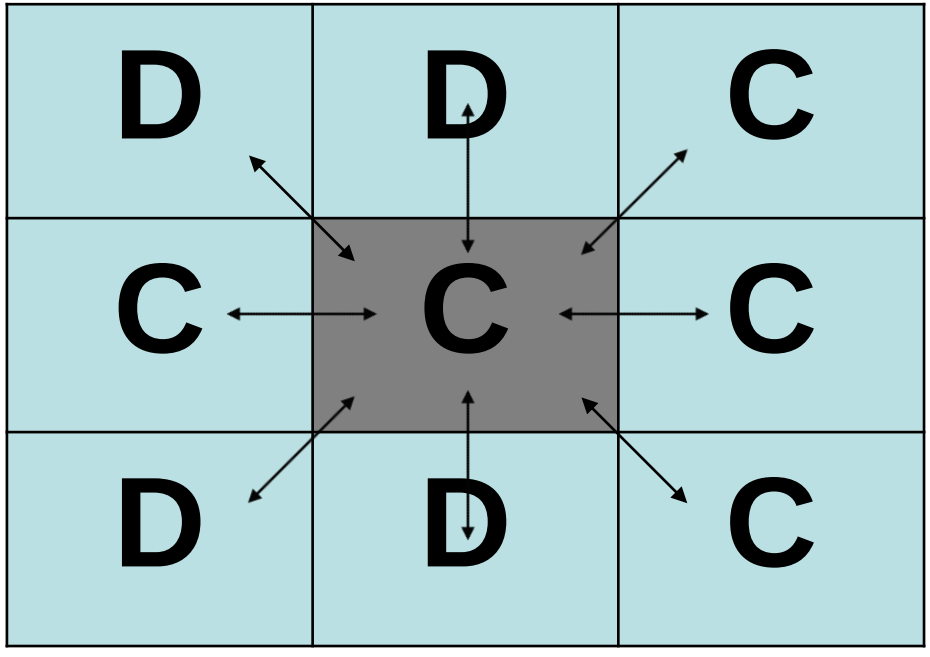
What is cooperation? Nowak's spatial game

Agents 1 and 2 cooperate on task T

$$Profit = Utility - Cost \qquad (or\ \Pi = U - C)$$

Agent 1 and 2 cooperate	$\Pi = U/2 - C/2$	for each
Agent 1 or 2 defect	$\Pi = U/2$	for defector
	$\Pi = U/2 - k \cdot C$	for cooperater
Agent 1 and 2 defect	$\Pi = 0$	for each

What is cooperation? Nowak's spatial game



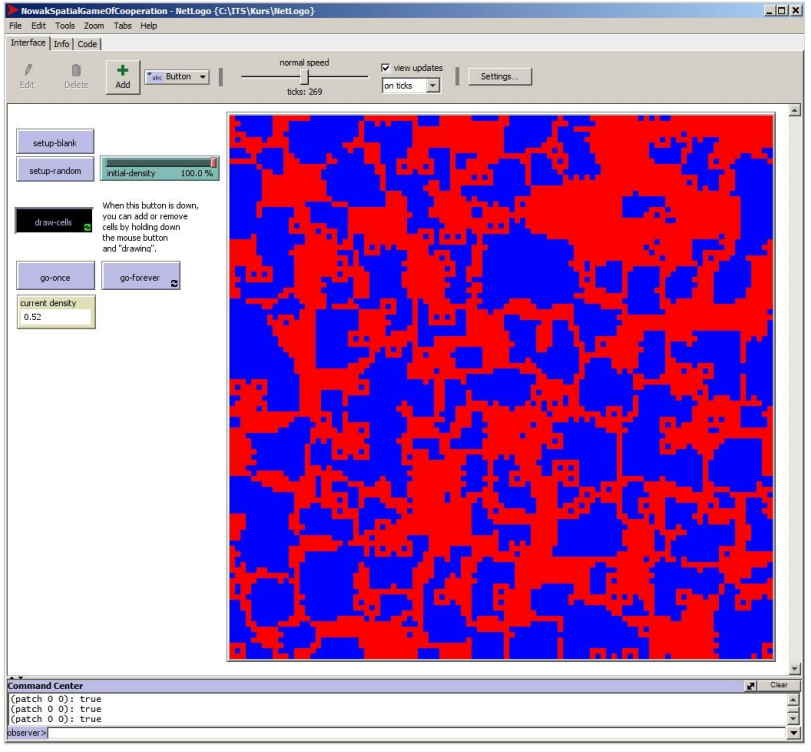
What is cooperation? Nowak's spatial game

NetLogo of NWU

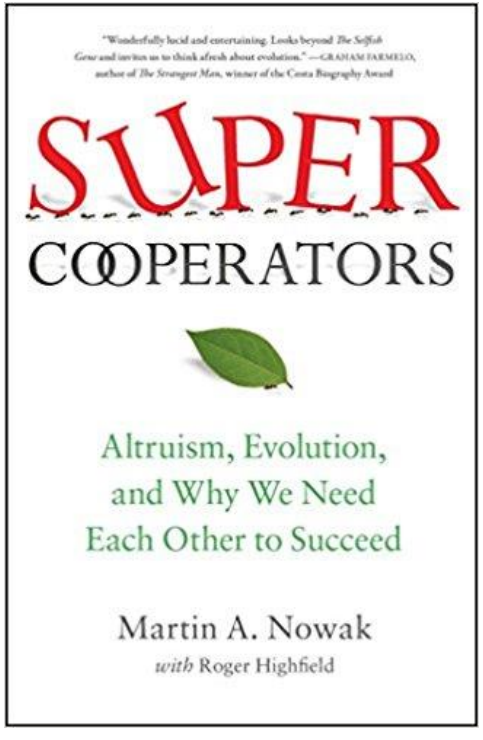
Framework for MAS

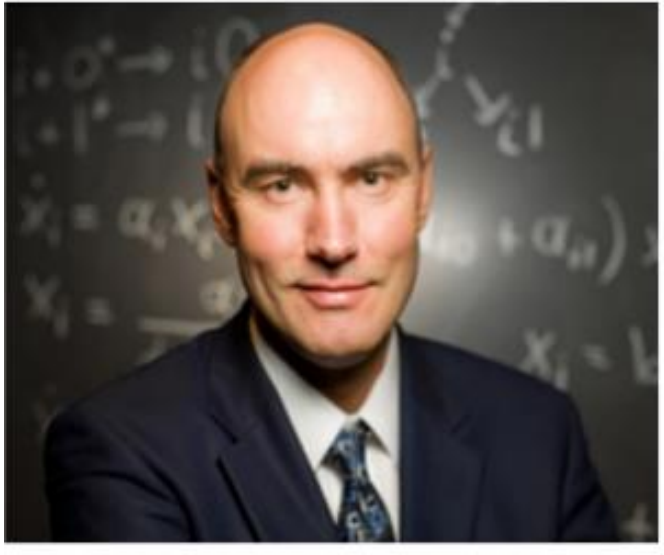
Blue = C (*Cooperator*)

Red = D (*Defector*)



What is cooperation? Nowak's spatial game





Professor Martin Nowak
Harvard University

Modern swam robotics

Some examples:

- FFI Valkyrie
- Harvard Kilobots and others

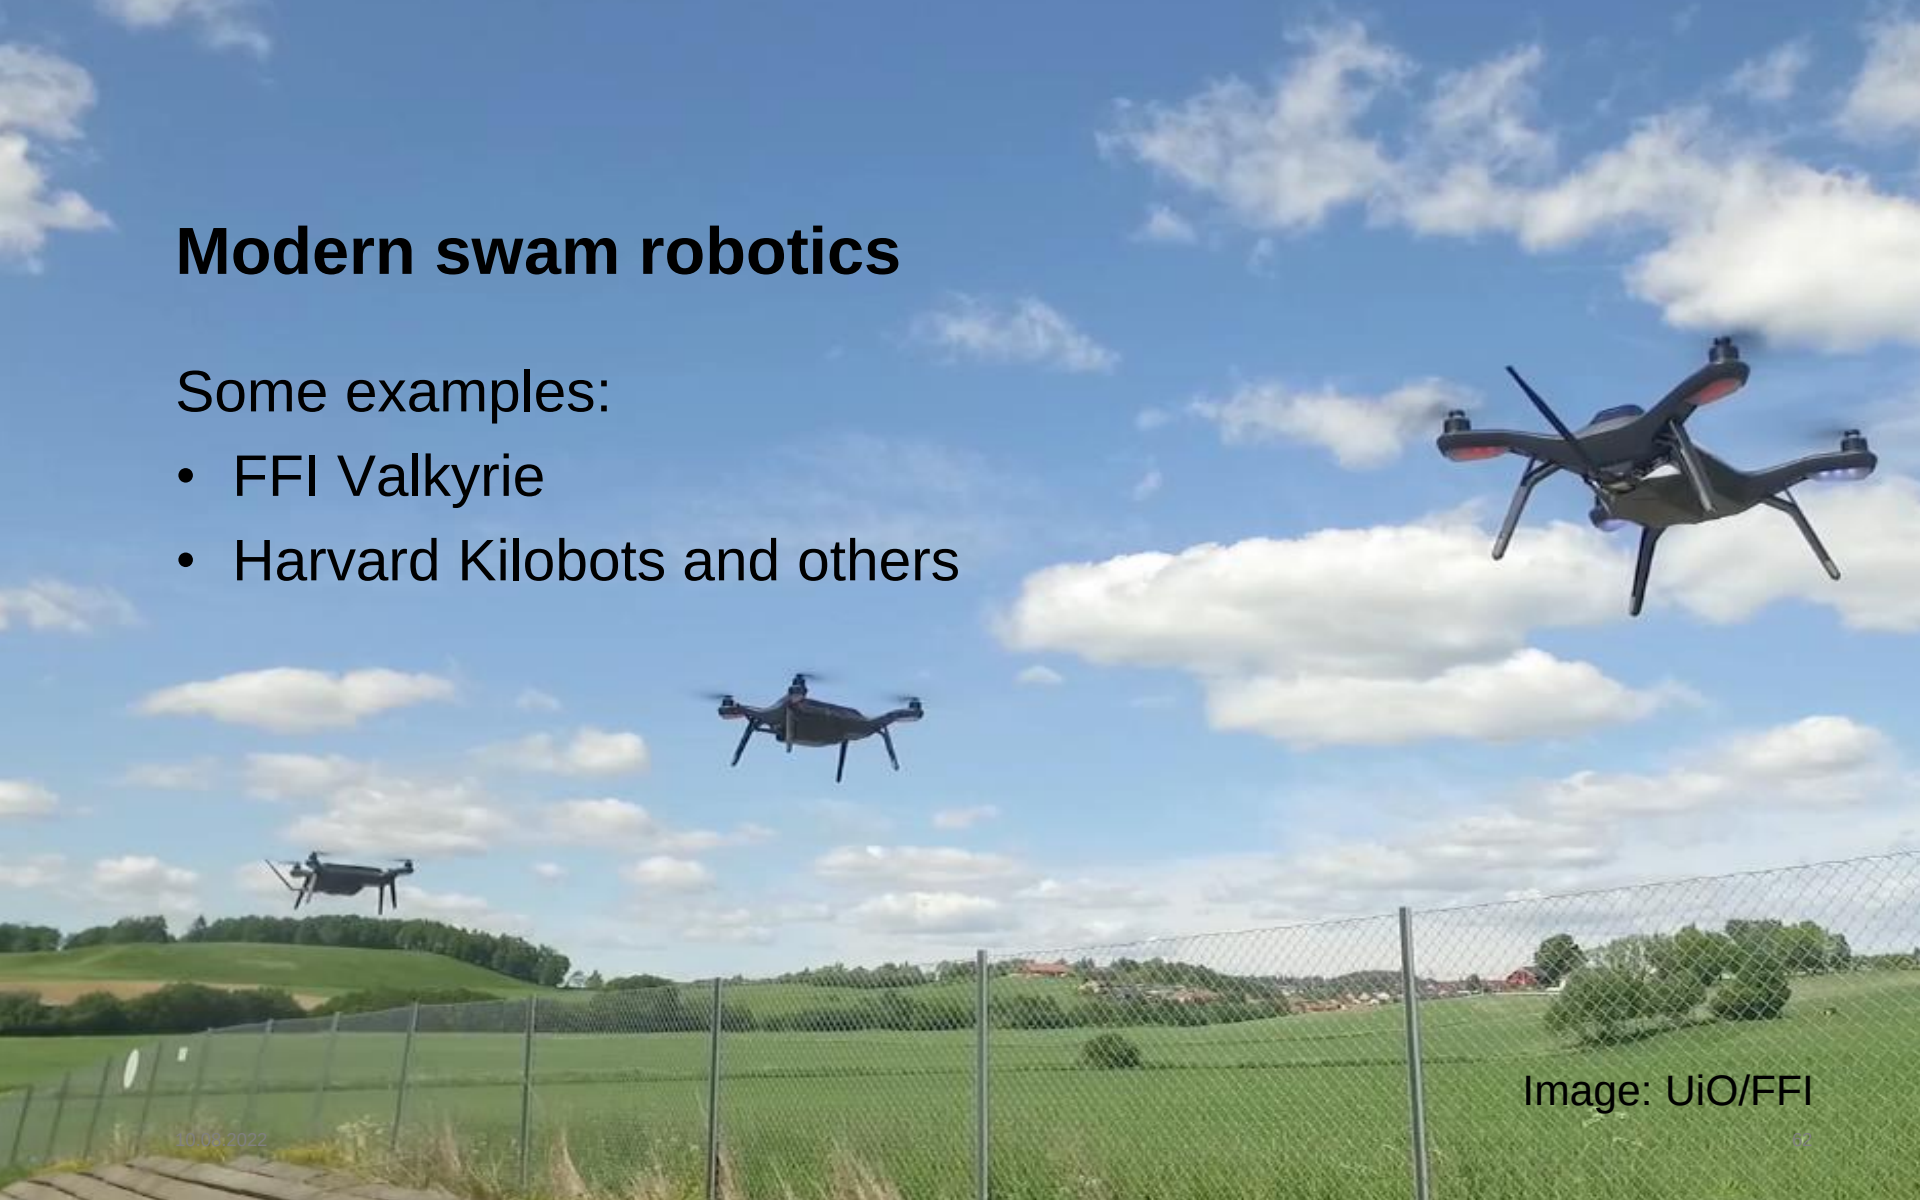


Image: UiO/FFI

Modern swam robotics: FFI Valkyrie

Why UAVs?

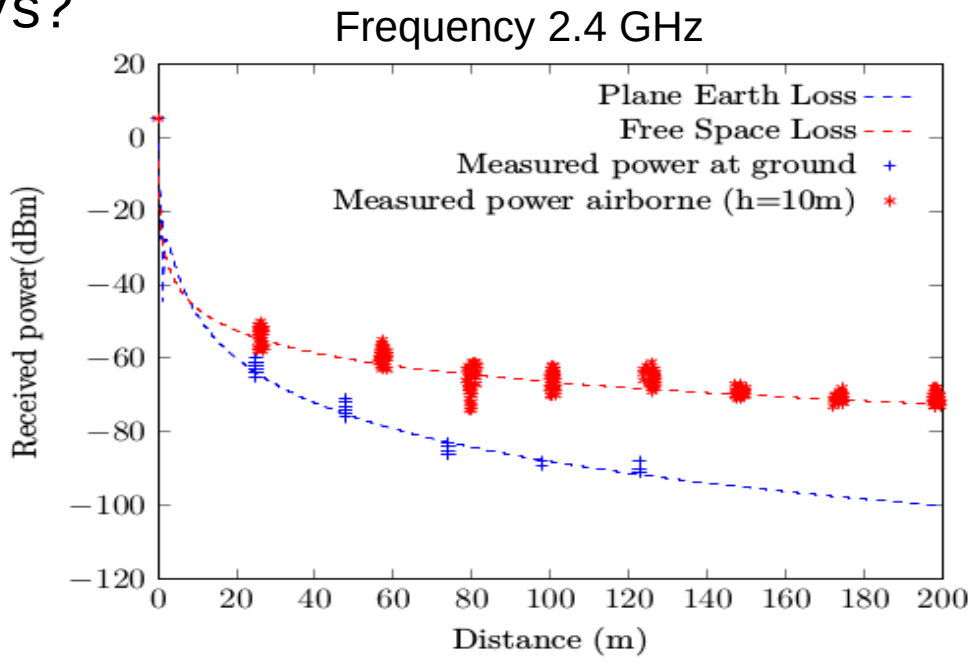
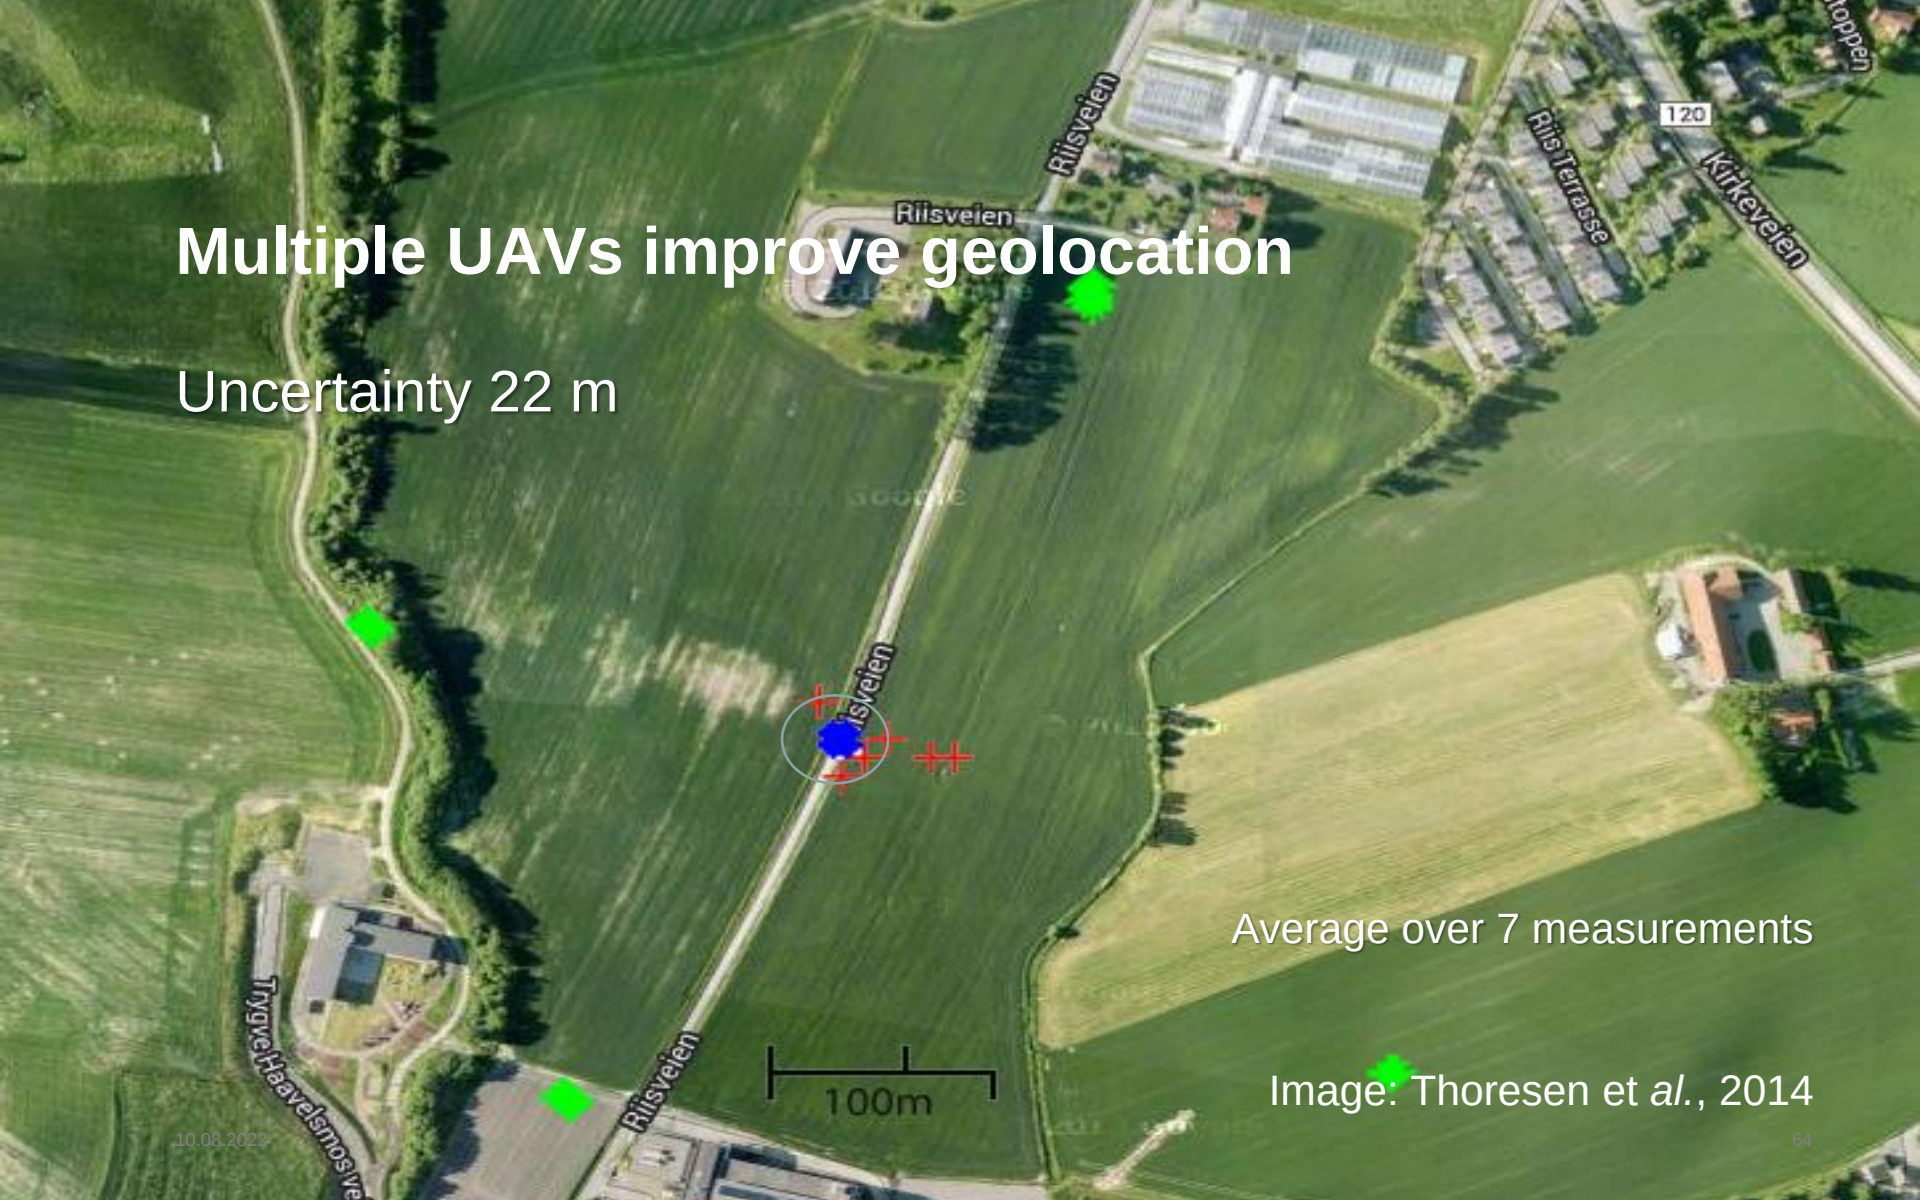


Image: Thoresen et *al.*, 2014

Multiple UAVs improve geolocation

Uncertainty 22 m



Average over 7 measurements

Image: Thoresen et *al.*, 2014

FFI Valkyrie



Autonomous communication and sensor network

Image: 2018 Sondre Engebråten, FFI

7 incredible modern swarm robots



Image: Harvard

Summary of lecture 1

Multiagent systems are about ‘computers’ working together in:

- Pure software systems
- Cyber-physical systems
- Robotic systems

and the research ranges from applied technology to fundamental science.