Large-scale adaptive systems

LECTURE 2: ADAPTATION MECHANISMS

Dr. Stefan Dulman
s.o.dulman@tudelft.nl

Review previous lecture

- Concepts
  - Self-Organization
  - Emergent Phenomena
  - Decentralized Control
  - Adaptation
  - Dynamic Change
  - Complexity

Adaptation

- Oxford dictionary:

  the process of change by which an organism or species becomes better suited to its environment

- Systems in nature provide *inspiration*
- *Inspiration* leads to engineered systems

Adaptable systems

- iGoogle, Blackboard, etc.
  - Customization of the way something (a page) looks

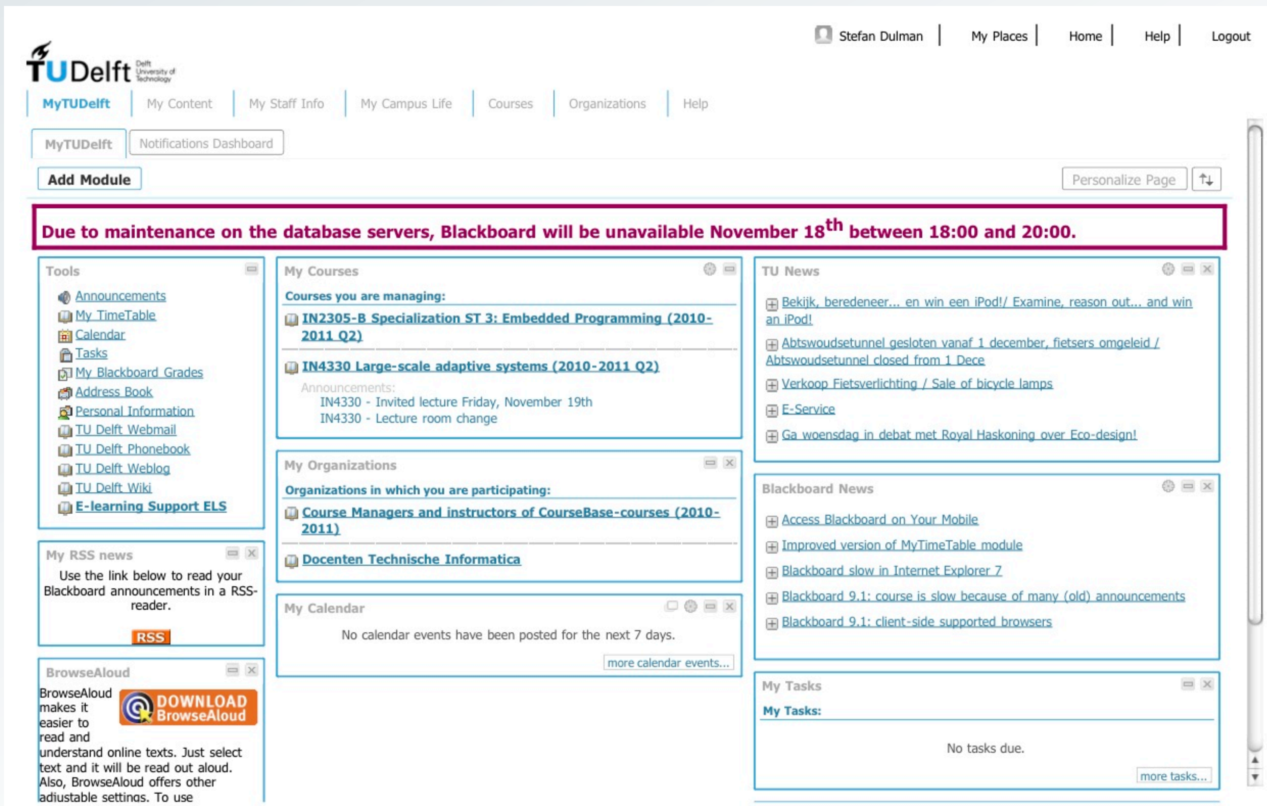

Terminology

- Self-Adaptive Systems - not in the dictionary …
  - "Self-adaptive systems work in a top down manner. They evaluate their own global behavior and change it when the evaluation indicates that they are not accomplishing what they were intended to do, or when better functionality or performance is possible.
  - Self-organizing systems work bottom up. They are composed of a large number of components that interact locally according to typically simple rules. The global behavior of the system emerges from these local interactions." (http://www.saso-conference.org/)

- Adapts to changing environment, on its own

Self-adaptive systems

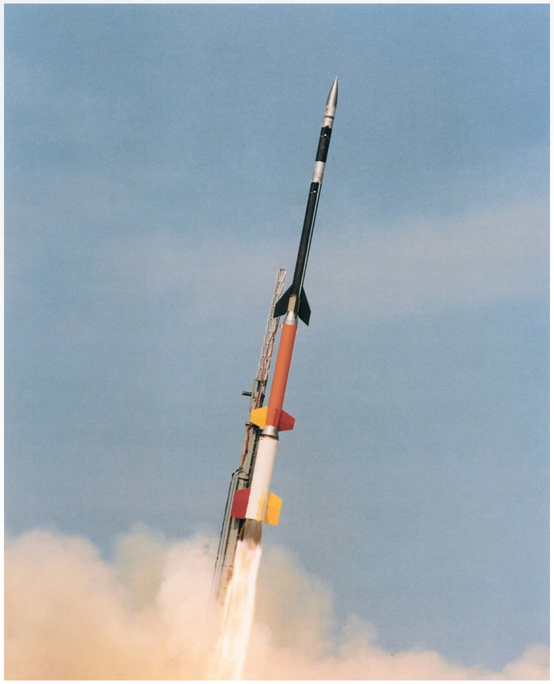

Self-organizing systems

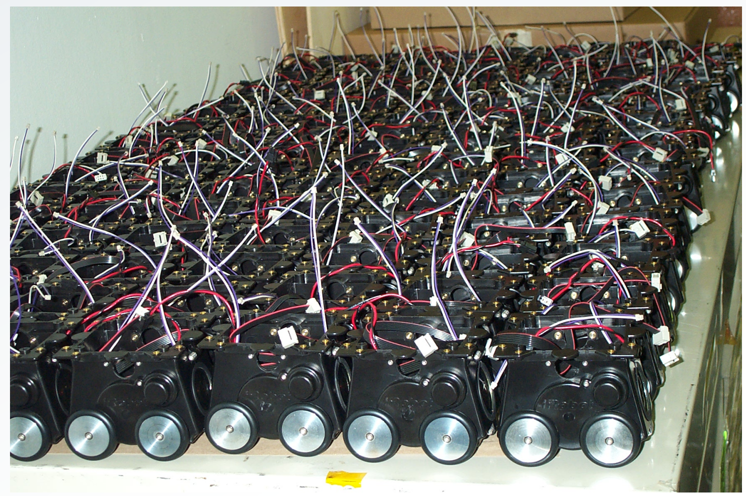

Lecture 2: Overview

- *Adaptation Mechanisms*
  - *Introduction*
  - Feedback mechanisms
    - Example: Schools of fish, flocks of birds
  - Stigmergy
    - Example: Economic based mechanisms
  - Autopoiesis
  - Reinforcement learning
    - Example: trust-based system

Feedback loops

- Adaptation is a response to feedback loops
- Positive feedback
  - Synonyms: self-enhancement, amplification, facilitation, autocatalysis
  - Amplification of fluctuations
- Negative feedback

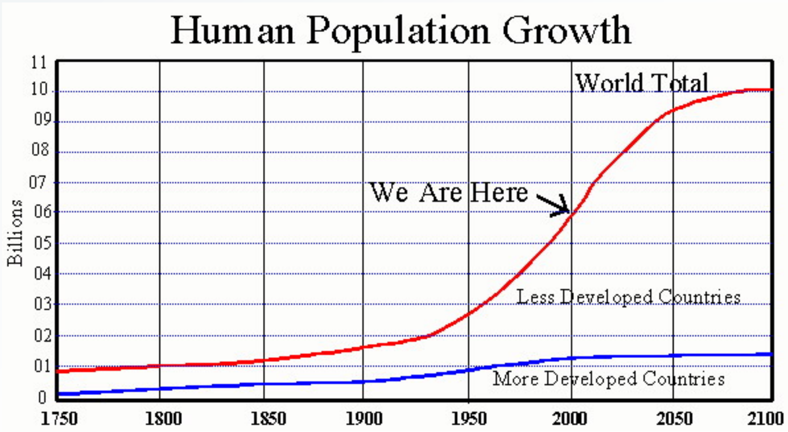

Positive feedback isn't always negative

M. Resnink – Learning about life

Control theory

- Open loop controller (feed-forward)
  - No relation between output (y) and input (u)
  - Controller responds to disturbance in a pre-defined way
    - controller does not compensate for unexpected changes
  - ′ Car speed is fixed ′
    - Car slows down when climbs a hill: no additional compensation from controller

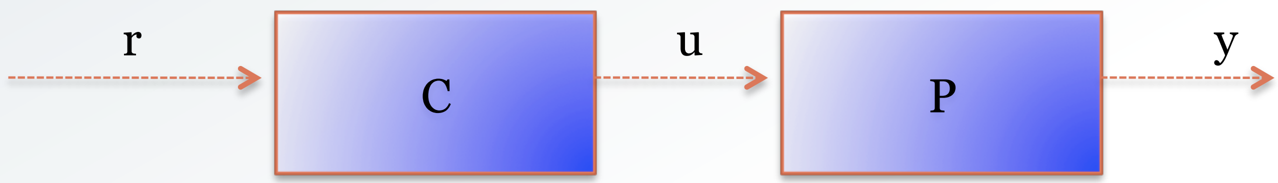

Control theory

- Closed loop controller (feedback)
  - Inputs (u) have an effect on the output (y)
  - Sensors monitor the output (y)
  - Controller monitors the error (e) and adjusts the inputs
    - Ex.: cruise control of vehicles

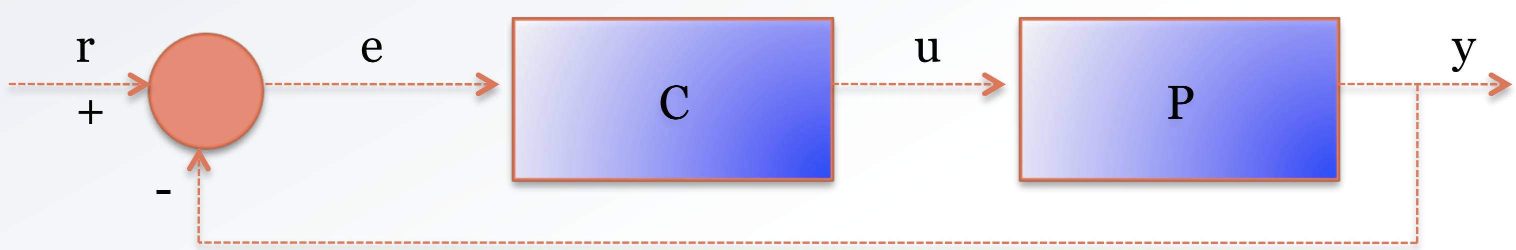

Control theory

- Ex.: PID controller
  - *Proportional* value: reaction to the current error
  - *Integral* value: reaction on the sum of recent errors
  - *Derivative*: reaction to the rate of change in errors
- "Intelligent" control
  - Incorporates AI computing techniques
  - Neural networks, fuzzy logic, machine learning evolutionary computation
- Stability, oscillation, self-stabilisation

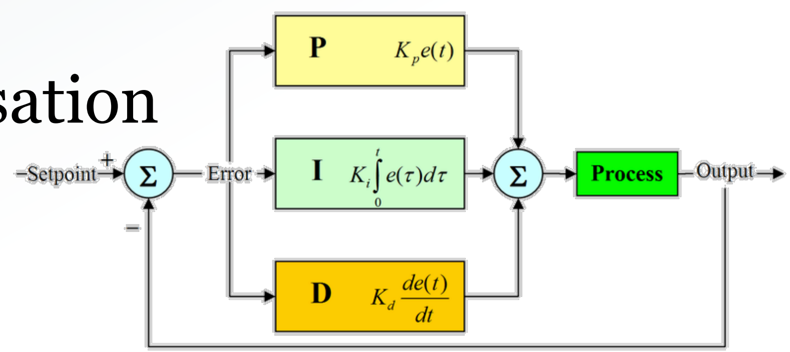

Feedback realization

- Information can be passed as:
  - Communication
  - Stigmergy – altering the environment
- Large-scale system in a *dynamic* environment
  - Combination of the above
  - Local actions – global response

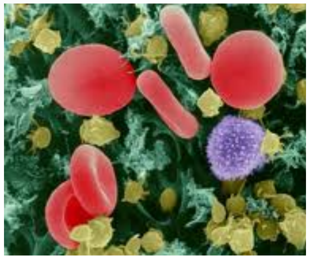

Communication-based feedback

- Direct communication among components of self-organising system
- Schooling and Flocking
  - Wave of reaction: communicated progressively to all components of school, or flock
  - Needed:
    - Monitoring of position and speed of neighbours
    - Adaptation of own position and speed
- Feedback mechanisms
  - Maintain a *minimum distance* from other objects in the environment, including other fishes/birds
  - *Move toward the perceived centre of mass* of fishes/birds in its neighbourhood
  - Match velocities with neighbours

Schools of Fish, Flocks of Birds

- Fish
  - Visual alignment: attraction effect / Direction
  - Sound receptors: Lateral line
    - Canals located at the side of fish act as receptors
    - Repulsion effect / Distance and Speed
- Birds
  - Wings and tails marking

Adaptation

- Collective Defence
  - Zigzagging activity, separating group + reforming
  - Act as a “Wall” against attacker
- Collective Feeding Activity
  - Encircling of a group of preys (tuna, whales)

Fish Manoeuvre to avoid Predators

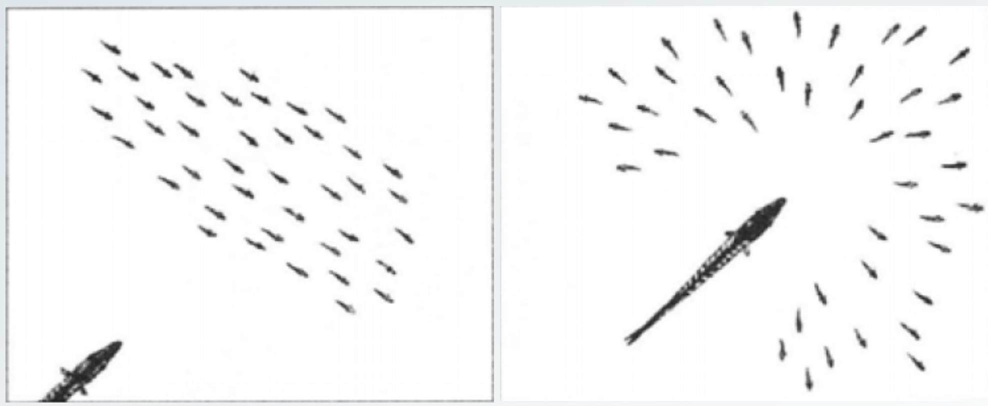

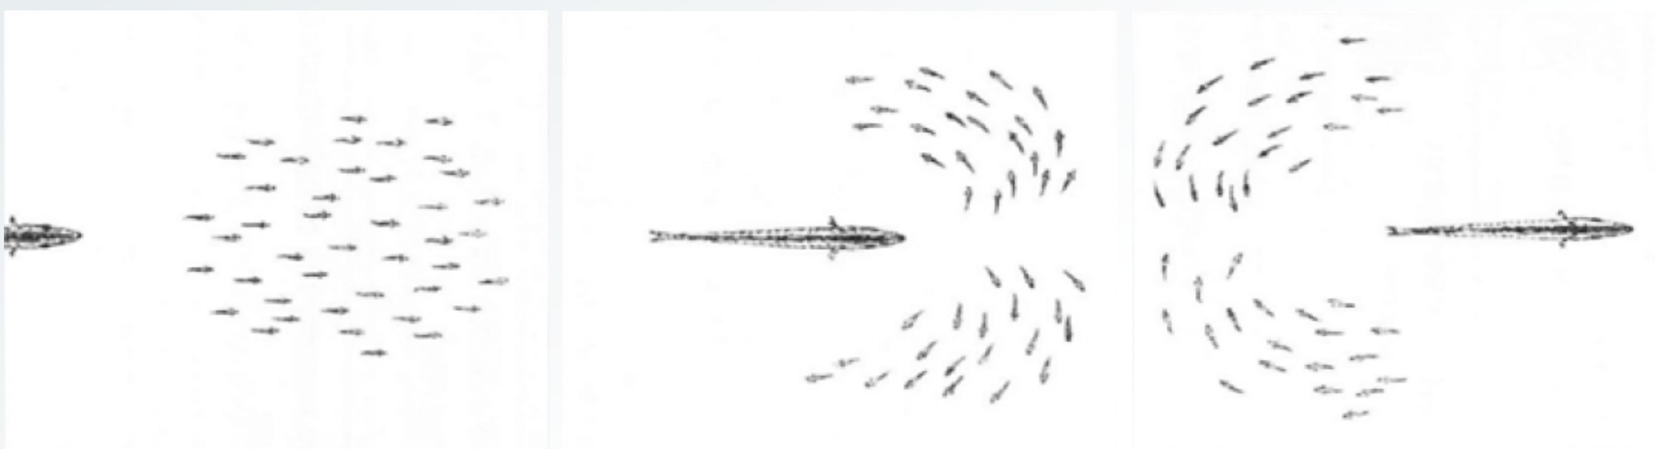

Boids

- Craig Reynolds
  - http://www.red3d.com/cwr/boids
- Reaction (apply rules) only based on the flock mates in a small neighbourhood
- Neighbourhood is defined by:
  - Distance
  - Angle
  - Taken from boid's direction
- Boids
  - Obstacle avoidance
  - Breaking the Ice

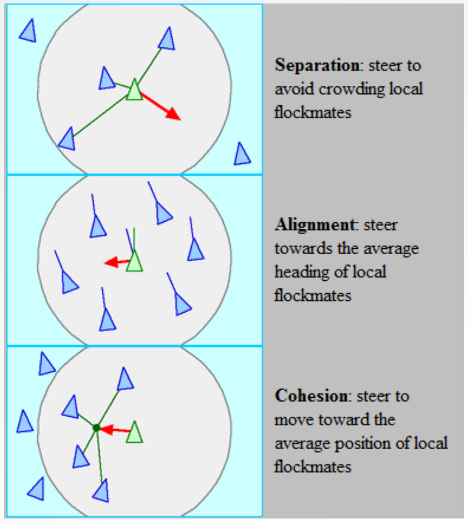

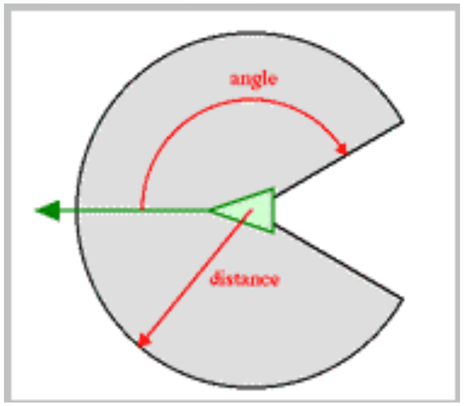

Lecture 2: Overview

- *Adaptation Mechanisms*
  - *Introduction*
  - *Feedback mechanisms*
    - *Example: Schools of fish, flocks of birds*
  - Stigmergy
    - Example: Economic based mechanisms
  - Autopoiesis
  - Reinforcement learning
    - Example: trust-based system

Stigmergy

- Introduction
  - Pierre-Paul Grassé, French biologist, 1959
    - Proposed theory of stigmergy while observing termites
    - Meaning: “incite to work”
- Definition
  - Indirect communication among components of a self-organising system
  - Mechanism: individual components modify their local environment
- Modification to environment
  - Pheromone (quantitative stigmergy)
    - E.g. foraging ants trails
  - Work-in-progress (qualitative stigmergy)
    - E.g. wasps nests construction

Stigmergy - mechanism

- Pheromones
  - Chemical volatile substance
  - Deposited into environment by individual
  - Retrieved (sensed) by individuals
  - Positive feedback: attractive effect
  - Example: ants' trails
- Pheromone description
  - Life time: 30-60 min; frequency: 5 marks/20 cm
  - Type of pheromone
    - alarm pheromones, food trail pheromones
- Pheromone concentration
  - Quantity of pheromone deposited
  - Flux of components (rate of ants)
  - Evaporation rate=Concentration/Life time

Qualitative Stigmergy

- Work-in-progress
  - Stimulus provided by previous work
    - Sequence of stimulating configurations
    - Local stimulating configurations
      - Different at each stage
  - Wasps nest construction
  - Different distinct phases:
    - Initial, first cell, other cells
    - Cells added in a particular way

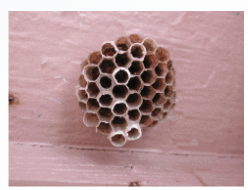

Work in progress

- Wasps apply rules to decide on what to do next
  - Start with one "brick"
  - Deposit new "bricks" depending on configuration
    - Bricks cannot be removed
    - Rule may be deterministic or probabilistic
- Rules:
  - Mapping: Configuration → Action (look-up table)
  - Non-conflicting (one rule for each configuration)
  - No overlap between different stages of building
  - Termination:
    - No more stimulating configuration
    - Separate rule based on the obtained size

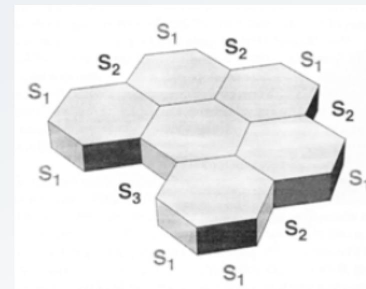

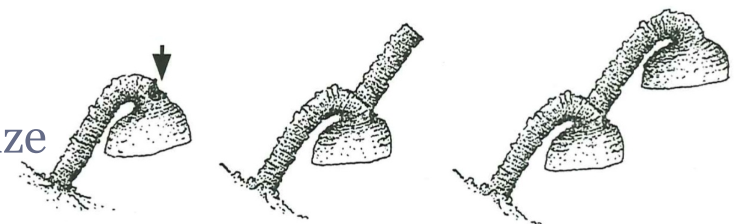

Work in progress

- Applications
  - Self-assembly of machines, of robots
  - Spatial application

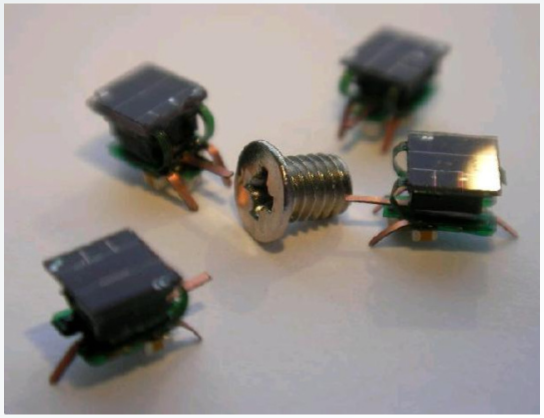

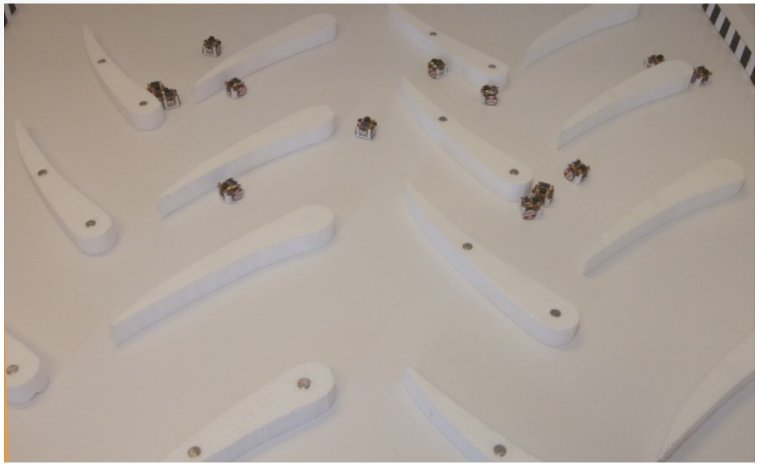

Examples - stigmergy

- WWW
  - A stigmergic communication medium for human
    - Everybody can upload (write) / download (read) information
- Wiki: Wikipedia
  - Initial user leaves an idea
  - Other users attracted by idea (add/modify content)
  - Result: complex structure of ideas/explanations/concepts
- Blogs
  - Communication through "boards"
  - Trails of information and links

Stigmergy - summary

- Ant-Pheromone Trails
  - Richest source of food
  - Shortest path: minimizing cost transport
  - One path (instead of two or more)
    - Strong path, no loss of ants, better defence against predators
- Termites
  - Adaptation of royal chamber to size of queen
  - Pillars distance, galleries size
- Wasps nests
  - Protections, defence
- Dynamic problem solving
  - Routing in telecommunication networks in dynamic environments

Check: http://en.wikipedia.org/wiki/Swarm_intelligence

Example: business mechanisms

- Based on dynamic business models and theories
  - Businesses are increasingly viewed as complex adaptive systems
  - Complex relationships between system components
  - Effect of changes in system or environment on system behaviour
- Personalised Marketing (one-to-one marketing)
  - Unique product offering for each customer
  - Individual offer to each customer
  - Differentiate a product from competing ones
- Phases
  - Identification of potential customers & their needs
  - Interaction with customers (learn about them)
  - Customisation of products, services, and communications

Adaptation of personalized marketing

- Personalised market strategy for each customer
  - evolves according to customer reactions
- Changing customers targeted
- Changing the prices quoted
  - Based on market dynamics
  - Based on customer characteristics
  - Based on the business goals

Example: Amazon

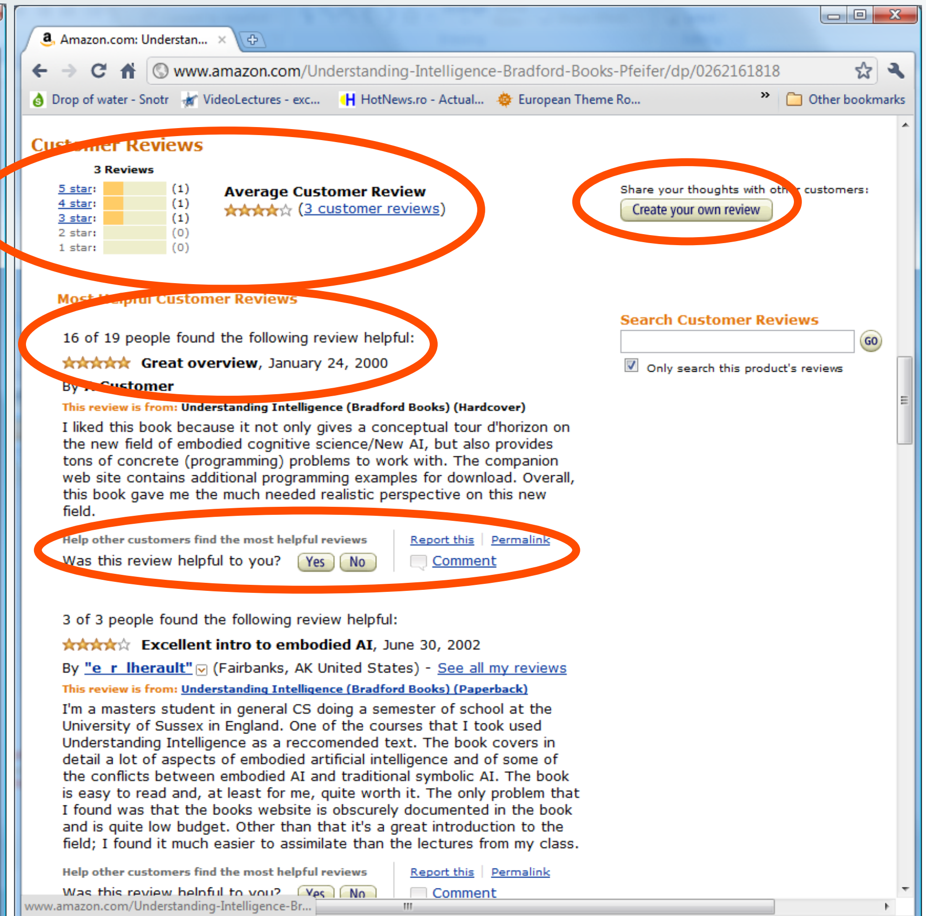

Lecture 2: Overview

- *Adaptation Mechanisms*
  - *Introduction*
  - *Feedback mechanisms*
    - *Example: Schools of fish, flocks of birds*
  - *Stigmergy*
    - *Example: Economic based mechanisms*
  - Autopoiesis
  - Reinforcement learning
    - Example: trust-based system

Autopoiesis

- Varela and Maturana 1971, Biologist
  - Study of living systems
  - Definition: the process through which an organisation is able to produce itself (self-production)
  - Applies to :
    - Systems made of autonomous components whose interactions self-maintain the system through the generation of system's components (cells, living organisms)
- Autopoietic systems (minimal living systems)
  - "Network of processes of production (synthesis and destruction) of components such that:
    - Components continuously regenerate and produce the network that produces them
    - Components constitute the system as a distinguishable unity in the domain in which they exist " (Varela 92)

Autopoiesis

- Varela Studies
  - Living systems, cognition, brain behaviour, consciousness
- Links and implications for:
  - Complex systems, brain studies, artificial life
- Adaptation
  - Change of components: interaction with environment

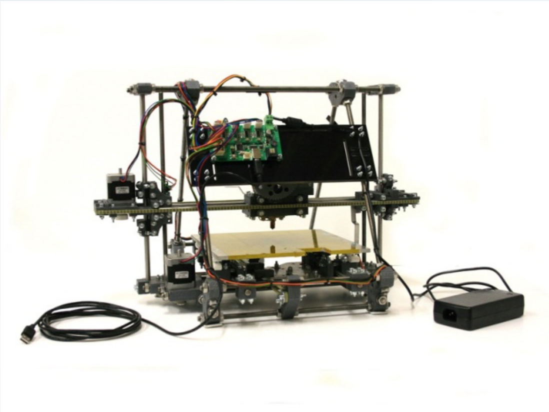

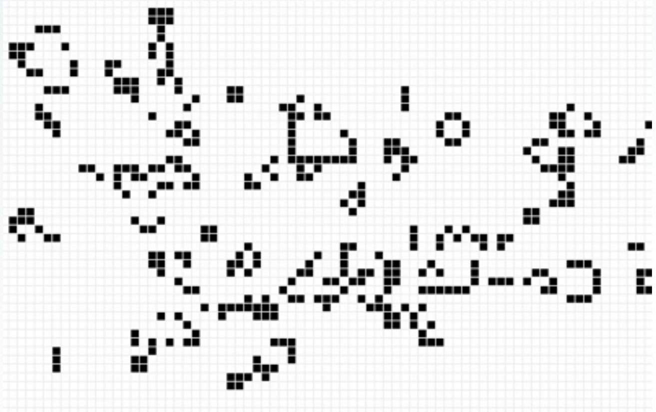

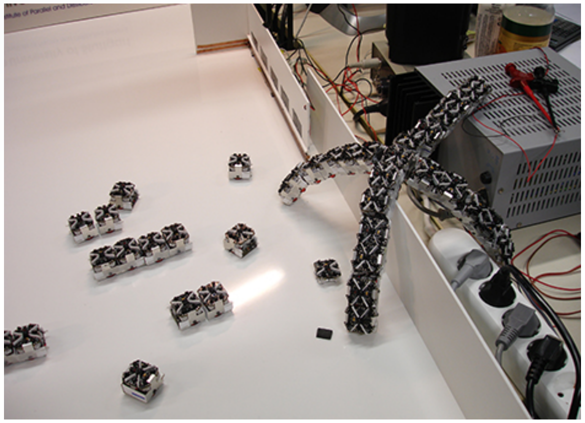

Lecture 2: Overview

- *Adaptation Mechanisms*
  - *Introduction*
  - *Feedback mechanisms*
    - *Example: Schools of fish, flocks of birds*
  - *Stigmergy*
    - *Example: Economic based mechanisms*
  - *Autopoiesis*
  - Reinforcement learning
    - Example: trust-based system

Reinforcement Learning

- Reinforcement learning (AI technique)
  - Agent learns behaviour through trial-and-error interactions with a dynamic uncertain environment
  - Programming agents by reward and punishment without needing to specify how the task is to be achieved
- Applies to:
  - Cases where it is difficult to determine what a program should do

I learned to ride with RL…

Reinforcement Learning

- Example
  - Program learning to ride bicycle
    - Maintain bicycle at 45° right and turn handle bars to the left
      - Fall down (45° right + turn left = bad)
    - Maintain bicycle at 45° right and turn handle bars to the right
      - Even worse (45° right + turn right or left = bad, 45° right = bad)
    - Maintain bicycle at 40° right and turn handle bars to the left
      - Bicycle goes to 45° right
    - Etc.
- Learning
  - State with immediate punishment must be avoided
  - State from which all actions lead to a state with immediate punishment must be avoided as well

Reinforcement Learning Model

- Reinforcement Learning Model
  - Environment's state must be observable (sensor inputs, mental representation, etc.)
  - Agent can observe perfectly well the system
- Elements:
  - Set of environment's state S = { s }
  - Set of agent's actions A = { a }
  - Reinforcement function:
    - Mapping: r: S x A → R (real numbers)
    - Set of scalar "rewards", "punishment", "nothing" in R (real numbers)
  - Goal = Reinforcement function:
    - Sum of future reinforcement the agents want to maximise

Reinforcement Learning problem

- Find a policy p : S → A
  - for maximizing cumulative reward for the agent over the course of the problem
- Difficulty
  - System is not told immediately if a specific action is good or bad
  - It is only when it gets the cumulative reward (at the end) that it knows if something was wrong
  - Difficult to know which of all the previous actions have to be avoided
  - Search space of behaviours to find the "best" one is infinite

Dynamic programming

- Two simple rules
  - Action that causes immediately a bad result
    - Do not do that action again when in the same state
    - Bicycle: turn handle bars to the right when bicycle at 45^o right (fall down immediately)
  - All actions possible from given state lead to bad results
    - Avoid to be in that state again
    - Bicycle: Avoid to be at 450 right
- Reinforcement limitations
  - Difficult to identify rewards and punishments (negative rewards)
  - Necessity to control all sources of reinforcement
  - Difficulty to create internal changes

Examples

- Games (Black Jack, …)
  - Black Jack
    - Win if sum of cards is <=21 but higher than dealer
  - At the end of each game a reward is provided
  - Computer learns on the basis of reward:
    - total value of cards
    - > 21 lost
  - Determine strategy through learning
    - E.g.
      - hit if (score<11)
stand otherwise

Trust-based Systems

- Principals: interacting set of entities (trusted or untrusted)
- Local trust values
- Evidence
  - Direct Observations: evaluated outcome of an interaction
  - Recommendations: asked or received (indirect observation)
- Scenario
  - Request or need of interaction
  - Decision: current trust value, evidence, risk implied by requested interaction
  - After interaction: trust value updated on the basis of evaluated outcome of the interaction
- Trust evolves with time, allows to adapt behaviour of principals

Example: Printers and Users

- Set of printers (not predefined)
- Set of users (using printers, not predefined)
- Knowledge of capabilities before interactions
  - Postscript/double-sided
- Memory of interactions outcome
  - Only single-sided, no printing
- Local trust value "computation" and "update"
- Propagation of recommendations
- Risks:
  - Losing time using a far located printer, printer runs out of paper, etc.

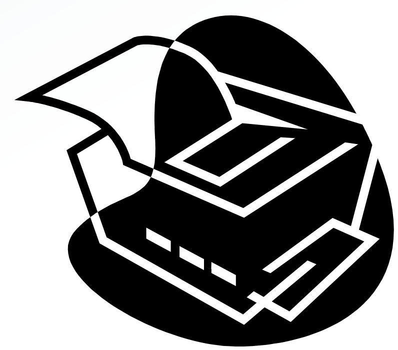

Printers and Users (1)

lw6: PostScript / Double-Sided/
Paper Jam / Problems with PDFs

lw3: New / Prints all PDFs

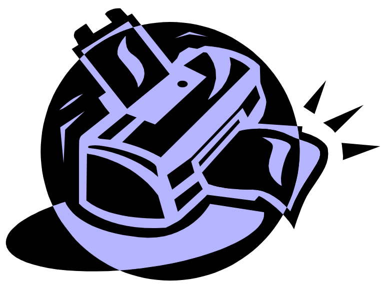

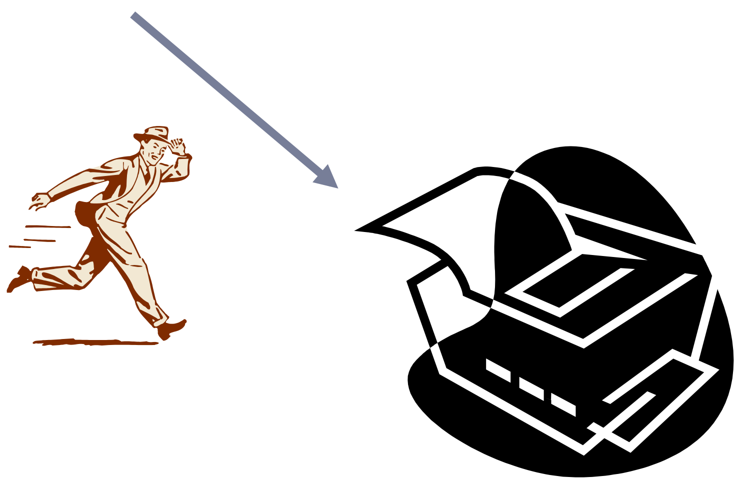

Printers and Users (2)

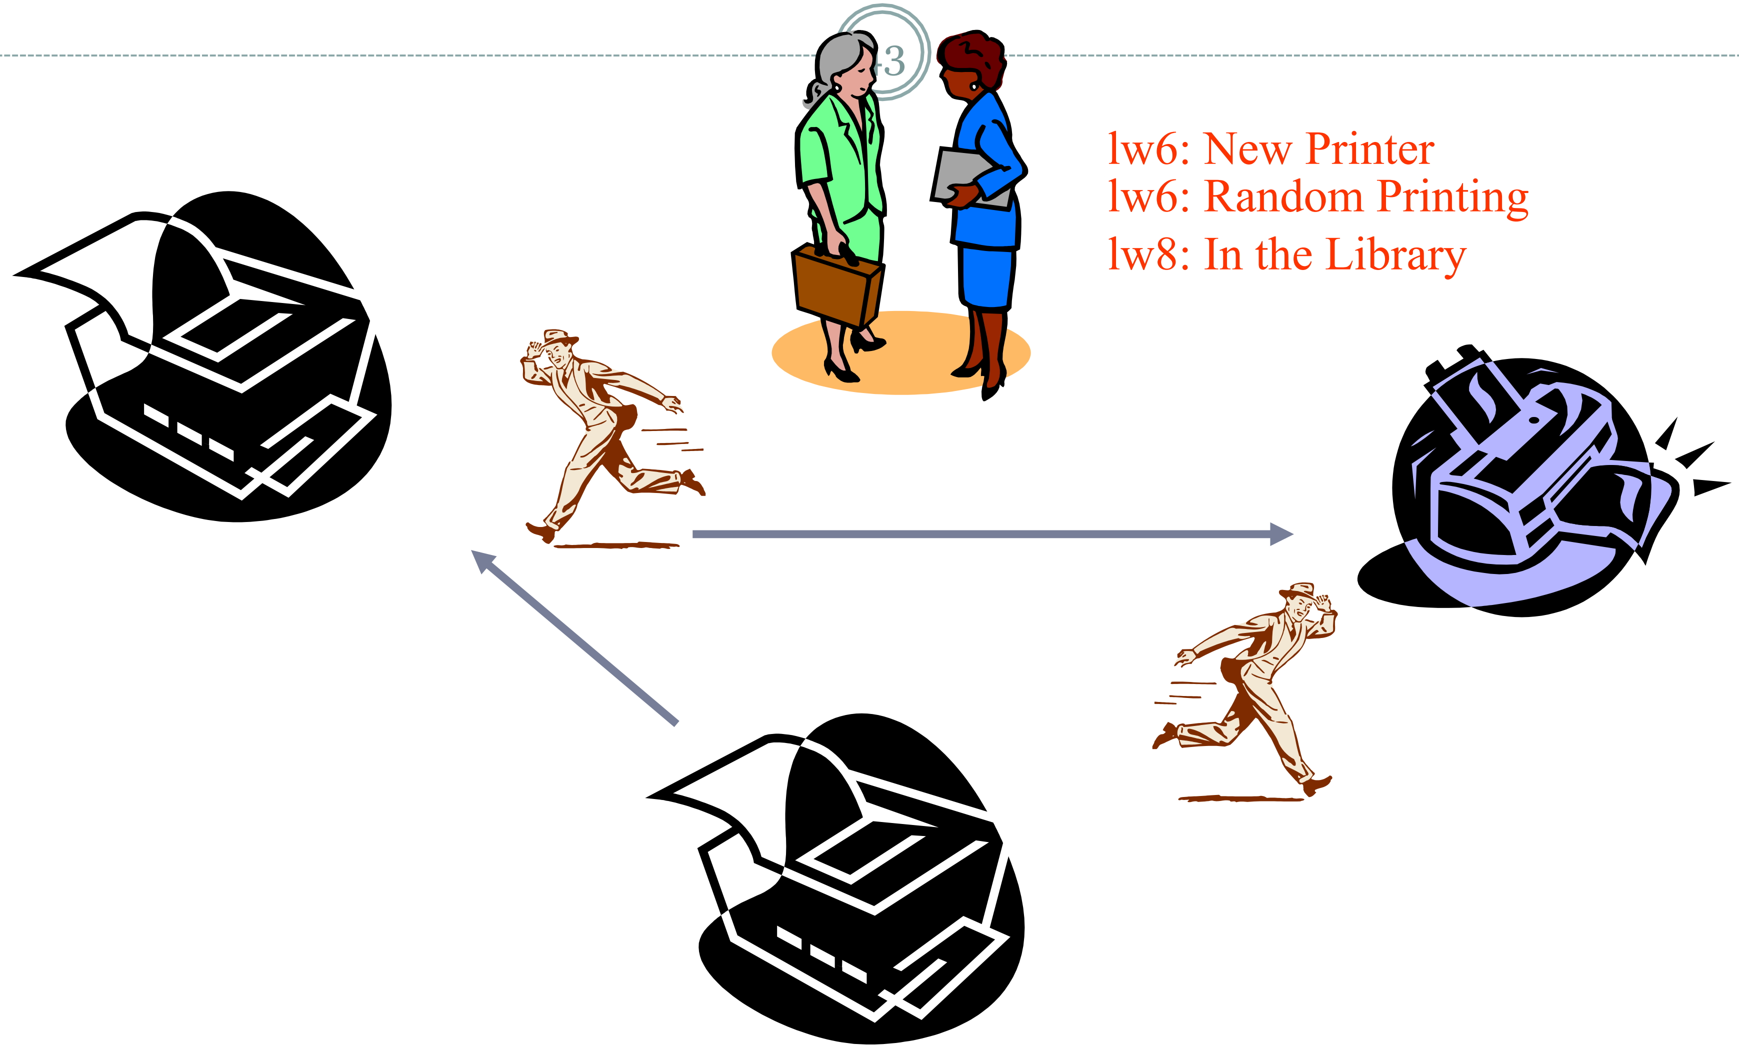

Printers and Users (3)

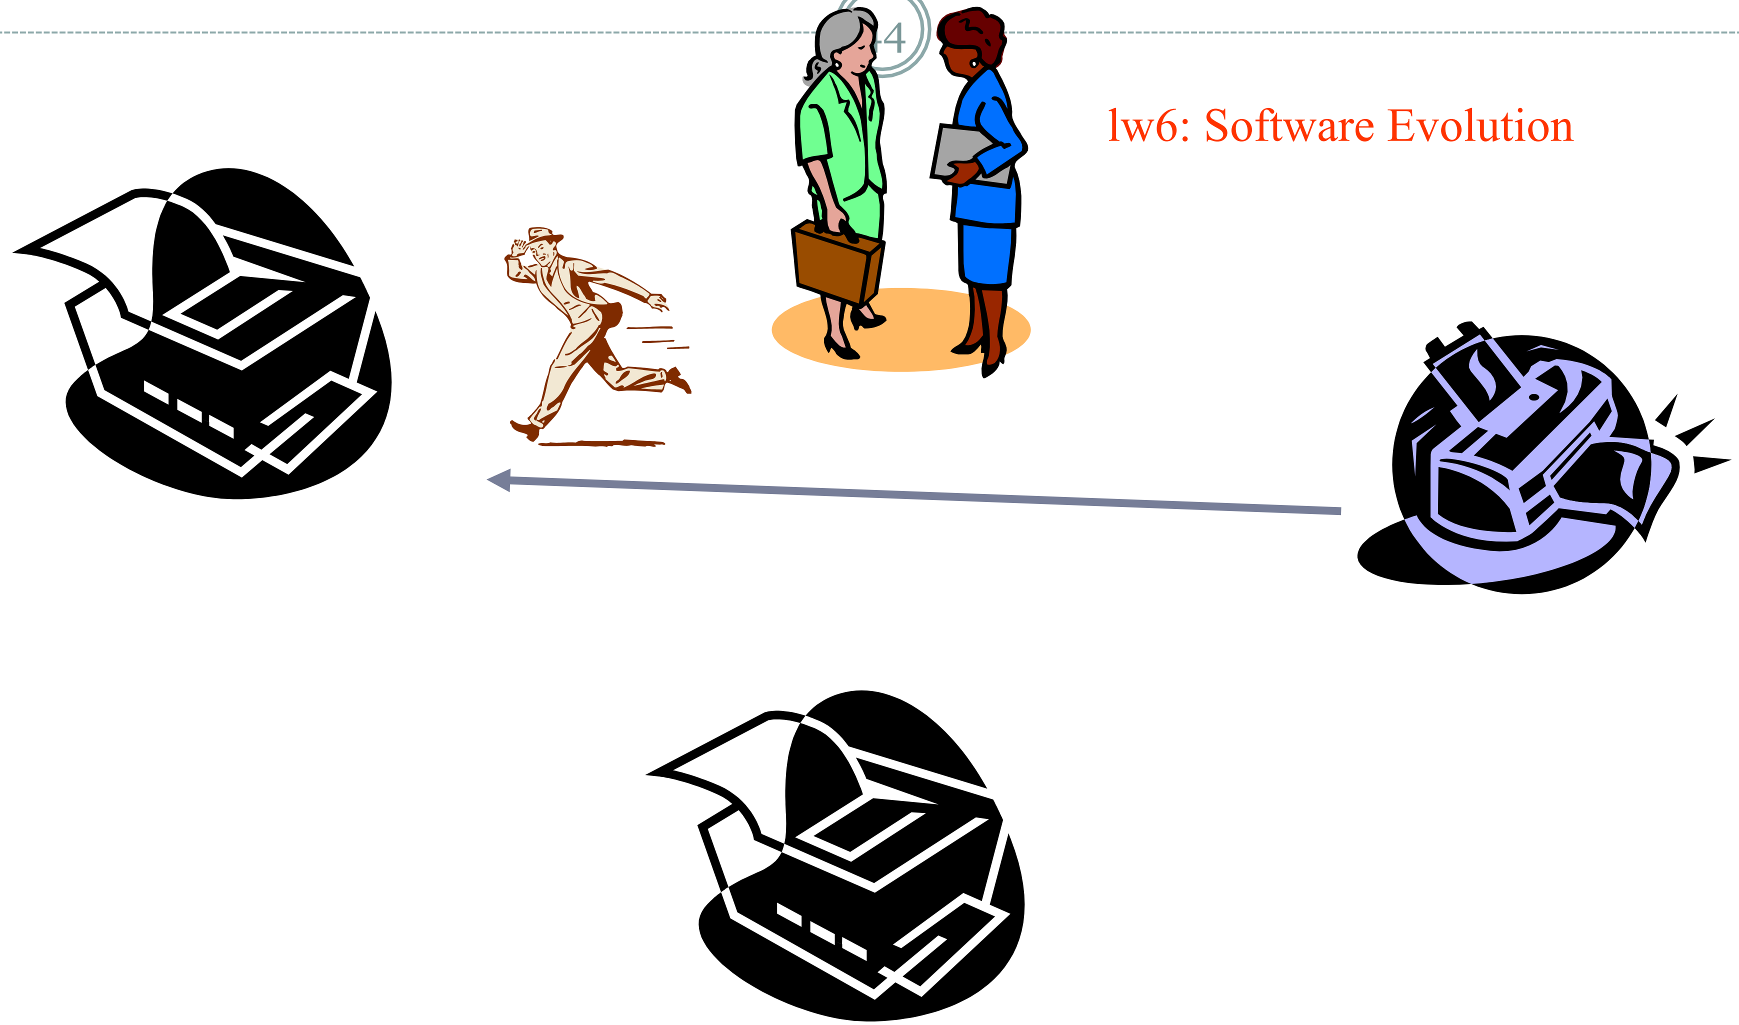

Trust as a SO mechanism (1)

- Mutual Causality
  - Exchange of recommendations
  - Direct observation
  - Users recommendations influence each other about the printer to use
- Autocatalysis
  - Positive evidence reinforces trust, and increases number of interactions
  - Negative evidence decreases trust, and decreases number of interactions
  - Trust in lw6 decreases, massive use of lw3

Trust as a SO mechanism (2)

- Far-from equilibrium condition
  - Power supply, network links, memory
  - Principals join and leave the system (autopoiesis)
  - Access denied to malfunctioning or malicious entities (entropy)
  - Faulty lw6 is no longer used
- Morphogenetic change
  - Random conditions affecting environment (broken network links)
  - Join/leave system
  - Software, hardware evolutions
  - lw6 updated two times (hardware, software)

Emergent Phenomena

- Reputation emerges from recommendations
  - lw6 is known to be unreliable
- Group formation emerges from interactions
  - Groups of users start/stop using printer
- Extension to artificial systems
  - Printers and PDAs
  - Printers maintain trust and reputation information about PDAs, possibly excluding them from printing

System functionality

- Task completion
  - Printing despite malfunctioning printers
- Optimise task
  - Print close to office
- More generally:
  - Modify own behaviour on the basis of current observed behaviour of neighbours or of interacting entities
  - Trade-off between risk (cost) and trust

Done!

- Adaptation Mechanisms
  - Introduction
  - Feedback mechanisms
    - Example: Schools of fish, flocks of birds
  - Stigmergy
    - Example: Economic based mechanisms
  - Autopoiesis
  - Reinforcement learning
    - Example: trust-based system