Special Topics in AI: Intelligent Agents and Multi-Agent Systems

Course Presentation and Introduction

Alessandro Farinelli

Outline

- Course Presentation
  - Aims, schedule, exam modalities
- Intelligent agents
  - AI, Intelligent agents, Rationality
- Multi-Agent Systems
  - Main features, techniques, applications

Lecture Material

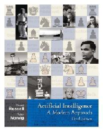

Artificial Intelligence – A Modern Approach

by Stuart Russell - Peter Norvig

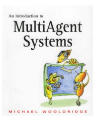

An Introduction to Multiagent Systems

by Michael Wooldridge

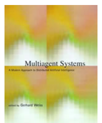

Multiagent Systems. 2nd Edition.

Gherard Weiss (Ed.)

Lecture slides and Info:

Course Organization

Wed 17th Oct. 15:30 -- 17:30 Room M;

Tue 23rd Oct. 15:30 -- 17:30; Room H

Tue 30th Oct. 15:30 -- 17:30; Room H

Mon. 5th Nov. 16:00 -- 18:00; Sala Verde

Tue. 13th Nov. 15:30 -- 17:30; Room H

Tue. 20th Nov. 15:30 -- 17:30; Room H

Tue. 27th Nov. 15:30 -- 17:30; Room H

Tue. 4th Dec. 15:30 -- 17:30; Room H

Tue. 11th Dec. 15:30 -- 17:30; Room H

Tue. 18th Dec. 15:30 -- 17:30; Room H

Course Aim

At the end of this course will be able to:

1. Understand main issues and research challenges for Multi-Agent Systems
   – Decentralized Coordination, Market Based Allocation, Reasoning under uncertainty
2. Model and solve Decentralized Coordination problems
   – DCOPs (exact and approx. methods)
3. Understand main models and solution techniques for decision making under uncertainty
   – MDP, POMDPs, Dec-MDPs

Course Program

1. Decentralized Coordination
   – Modeling Decentralized Coordination as DCOPs
   – DCOPs solution techniques (exact and approx.)
2. Market Based Allocation
   – Auction Mechanisms, Combinatorial auctions, Sequential auctions
3. Reasoning under uncertainty
   – MDPs, POMDPs
   – Probabilistic approaches for robot navigation

Exam modalities

- Students read, present to the class, and discuss a set of selected papers.
- Student together with instructor choose papers
  – Topics: Decentralized optimization, Market-Based Allocation, Reasoning under uncertainty (robotics)
- Presentation:
  – From 45mins to 1 hour + questions
  – During the last three lessons (4th 11th 18th Dec.)

Outline

- Course Presentation
  – Aims, schedule, exam modalities
- Intelligent agents
  – AI, Intelligent agents, Rationality
- Multi-Agent Systems
  – Main features, techniques, applications

What is AI?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Acting Humanly:Turing Test

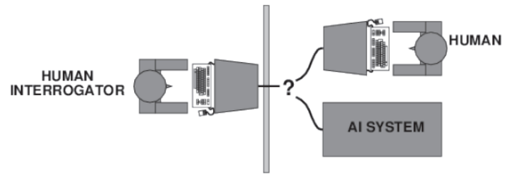

Turing(1950) Computing Machinery and Intelligence

- Can machine think? → Can machine behave like humans?
- Operational test: the imitation game

Problem: not reproducible, constructive or amenable to mathematical analysis

Thinking humanly: Cognitive Science

- Cognitive Neuroscience → theories of internal activities of the brains
  - Level of abstraction? Validation ?
- Available theories do not explain human-level intelligence

Thinking rationally: Laws of thoughts

- Normative not descriptive
- Problems:
  - Intelligence not always based on logical deliberation
  - What are the purpose of thinking ? Which thoughts should I have out of all the ones that I could have

Acting rationally

- Do the <u>right</u> thing
  - Action that maximizes some measure of performances given current information
- Thinking should be in service of rational actions
  - Thinking is not necessary (e.g., blinking reflex)
- Correct thinking (inference) does not always result in rational actions
  - Thinking is not sufficient

Rational agents

- Agent: entity that perceives and acts
- Rational agent
  - A function from percept histories to actions

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

  - For a given class of environments and tasks we seek the agent with best performance (<u>optimization problem</u>)

Agents and Environments

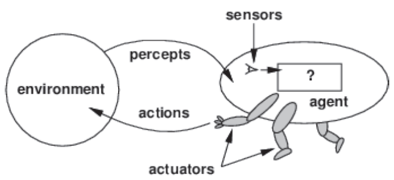

- <u>Agents</u>: humans, softbots, thermostats, robots, etc.
- <u>Agent function:</u> maps perception histories to actions
- <u>Agent program:</u> implements the agent function on the physical architecture

Rationality

- <u>Given</u> a performance measure for environment sequences
- <u>Rational agent:</u> chooses actions that maximizes the expected value given percept sequence
- Rational ≠ omniscient
  - Perception may not supply all relevant info
- Rational ≠ clairvoyant
  - Action outcome might be unexpected
- Hence Rational ≠ successful
- Rational => exploration, learning, autonomy,...
-

Agent Types: Simple reflex Agent

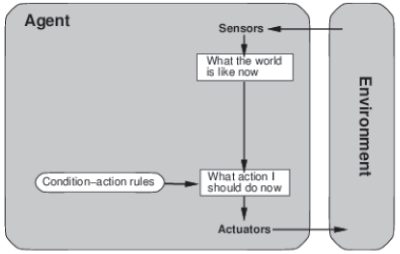

Agent Types: Goal-Based agents

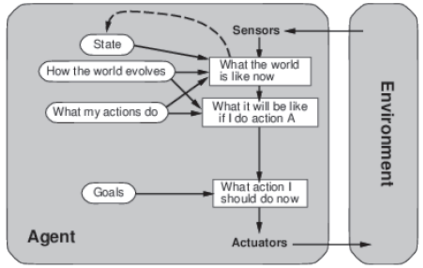

Agent Types: Utility-Based Agent

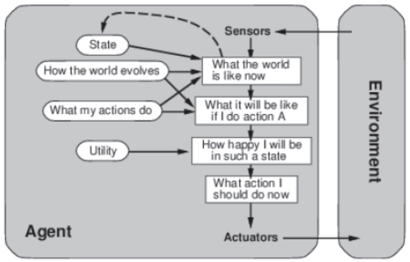

AI (recent) history

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's "Computing Machinery and Intelligence"
1950s	Early AI programs, e.g., Samuel's checkers program, Newell & Simon's Logic Theorist
1956	Dartmouth meeting: "Artificial Intelligence" adopted
1965	Robinson's complete algorithm for logical reasoning
1966–74	AI discovers computational complexity Neural network research almost disappears
1969–79	Early development of knowledge-based systems
1980–88	Expert systems industry booms
1988–93	Expert systems industry busts: "AI Winter"
1985–95	Neural networks return to popularity
1988–	Resurgence of probability; "Nouvelle AI": ALife, GAs, soft computing
1995–	Agents, agents, everywhere . . .
2003–	Human-level AI back on the agenda

AI Exciting Applications

- Game Playing
  - IBM’s Deep Blue (1997)
  - Poker (Now) http://webdocs.cs.ualberta.ca/~games/poker/
- Autonomous Control
  - Google self driving car http://www.ted.com/talks/sebastian_thrun_google_s_driverless_car.html
- Search and Recue/hostile environments
  - RoboCup Rescue (http://www.robocuprescue.org/)
- Human Agent Collectives
  - Orchid project (http://www.orchid.ac.uk/project-aims/)

Example: Search and Rescue

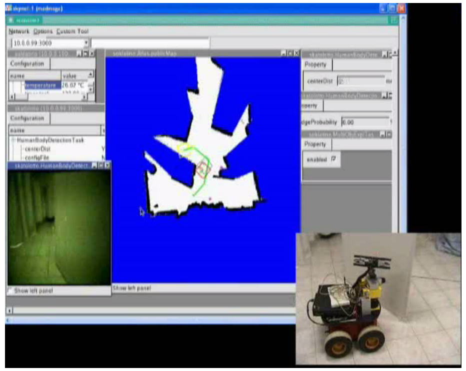

LabRoCoCo http://labrococo.dis.uniroma1.it/wiki/doku.php

Outline

- Course Presentation
  - Aims, schedule, exam modalities
- Intelligent agents
  - AI, Intelligent agents, Rationality
- Multi-Agent Systems
  - Main features, techniques, applications

Intelligent Agents

- Intelligent Agents: rational agent +
  - Reactivity
  - Pro-activeness
  - **Social ability** → Multi-Agent systems

Rational Agent

Intelligent Agent

Multi-Agent Systems

- (Durfee and Lesser 1989): “loosely coupled **network** of problem solvers that **interact** to solve problems that are **beyond** the **individual capabilities** or knowledge of each problem solver “
- Problem solvers: Intelligent agents
- (John Gage, Sun Microsystems)
  “The network is the computer”

MAS Characteristics

(K. P. Sycara 1998)

1. Each agent has incomplete information or capabilities for solving the problem and, thus, has a limited viewpoint
2. There is no system global control
3. Data is decentralized
4. Computation is asynchronous

Example: cooperative foraging

Why MAS?

- To solve problem that are too large for a single agent
  - Problem decomposition
- To Avoid single point of failure in critical applications
  - Disaster mitigation/urban search and rescue
- To model problem that are naturally described with collectives of autonomous components
  - Meeting scheduling, Traffic control, Forming coalition of customers, …

Main Research Areas in MAS

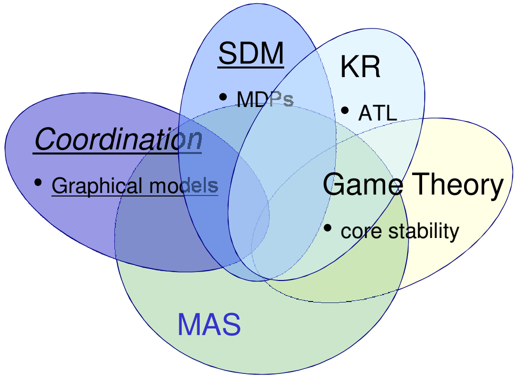

Applications of MAS I:

Games, entertainment and education

Real Time Strategy (e.g. Starcraft, Age of Empires)

→ group formation, task assignment, strategic planning

First Person Shooter (e.g. Half Life 2, Splinter Cell)

→character interactions

Applications of MAS II:

Search and Rescue

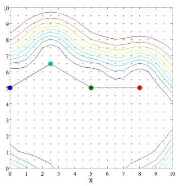

Cooperative information gathering

UAVs cooperative image collection

Cooperative information gathering

Limited communication

Joint work with Stranders, Rogers, Jennings [IJCAI 09]

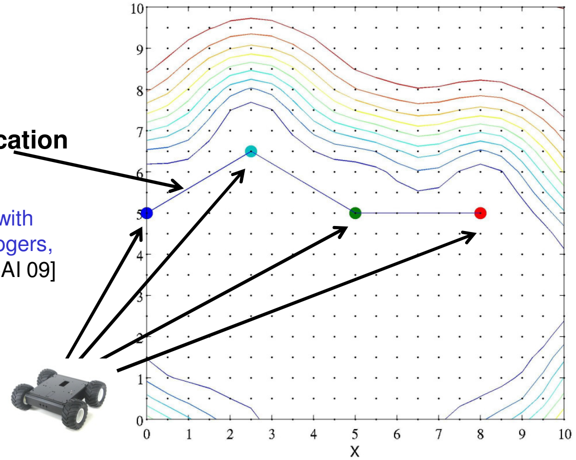

CIG: the model

- Monitor a spatial phenomena
- Model: scalar field
  - Two spatial dimensions
  - One temporal dimension

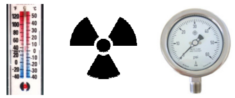

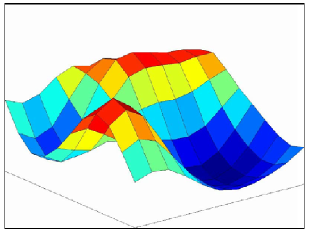

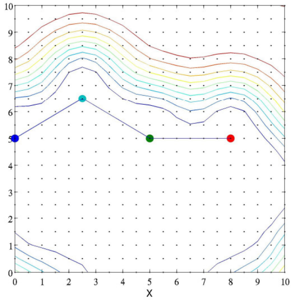

CIG: goal

- Minimise prediction uncertainty
- Given a measure here what is my uncertainty over there
- Tools:
  - Gaussian process
    - Estimate uncertainty
  - Entropy
    - Measure information

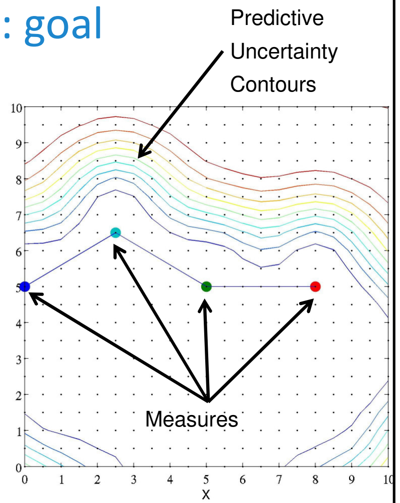

CIG: Performance measure and interactions

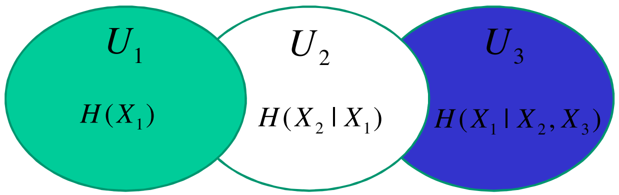

$$H(X_1, X_2, X_3) = H(X_1) + H(X_2 | X_1) + H(X_3 | X_1, X_2)$$

$$U_1 \quad + \quad U_2 \quad + \quad U_3 \quad = \sum U_i$$

CIG: Demo

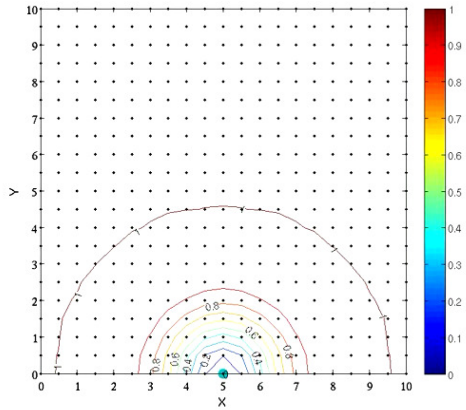

Cooperative Image Collection

Task Assignment for UAVs

Joint work with:
Delle Fave, Rogers,
Jennings

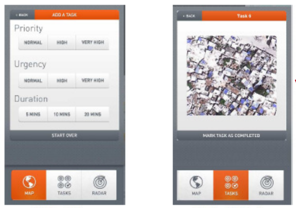

(a) Task Input (b) Task Monitor

Video Streaming

Interest points

Coordination

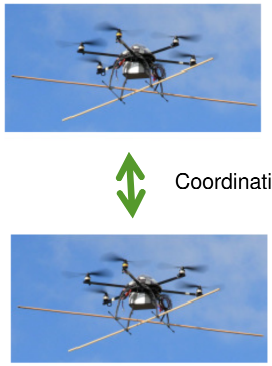

CIC: Task utility

Task completion

Priority

$$U_j(X_j) = p_j \cdot u_j^{t - t_j^0} \cdot [\,1 - e^{-\lambda_j \cdot (t_2 - t_1)}\,]$$

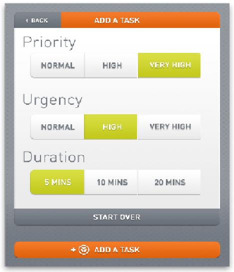

Urgency

First assigned UAVs reaches task

Last assigned UAVs leaves task (consider battery life)

38

CIC: Interactions

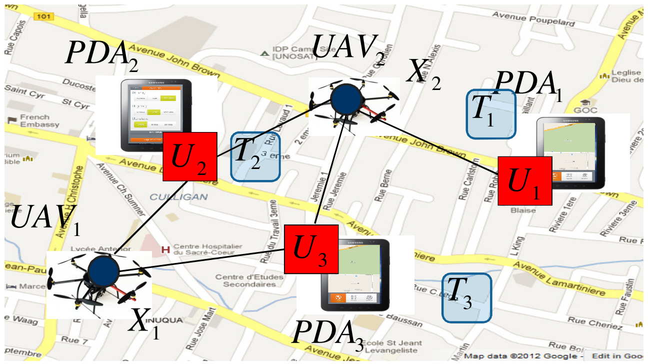

CIC: UAVs Demo

Applications of MAS III:
Energy management

Intelligent agents for the smart grid

Mechanism design and Energy trading
- Force demand to follow supply

Home Energy management
- Agents to decide load scheduling and storage

Collective energy trading
- Buy and sell energy as collectives

Electricity markets

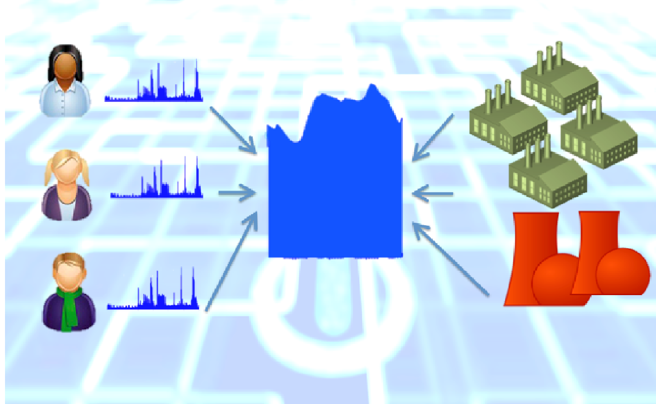

Electricity markets

Baseload: Carried by baseload stations with low cost generation, efficiency and safety

Electricity markets

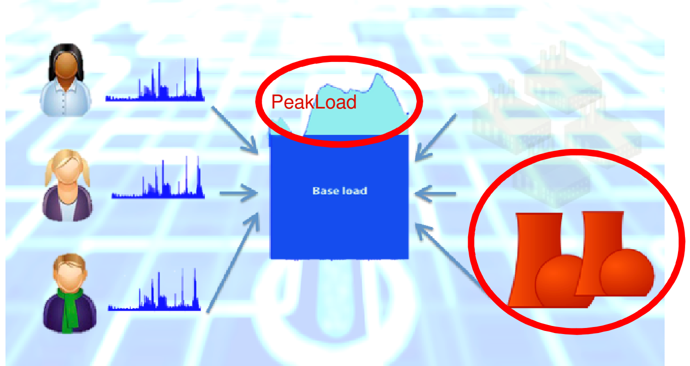

Peakload: Carried by expensive, carbon-intensive peaking plants generators

Electricity group purchasing

- Allow **group purchasing among electricity consumers**
- Very popular successful cases
  - Groupon, Groupalia
  - UK Labour party initiative on collective electricity purchase

Electricity Group Purchasing

- **Virtual Electricity Consumer (VEC):** A group of consumers that act in the market as a single energy consumer.

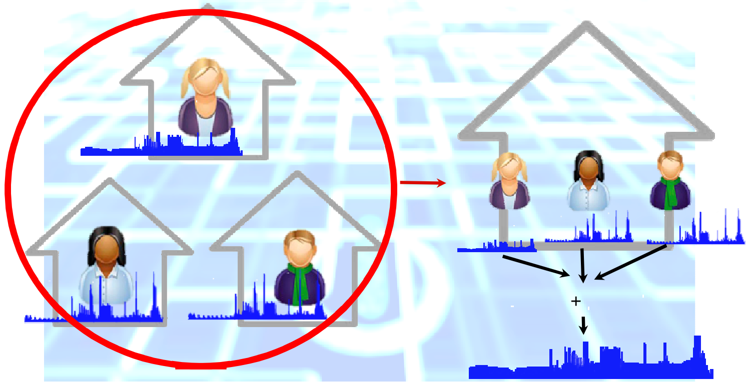

Group synergies

- Traditional group purchasing based on group size
- Group synergy: **complementary energy restrictions**
  - Flattened demand => Better prices

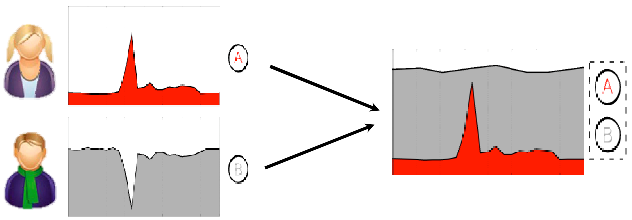

Social networks

- **Social networks** to support the VEC **formation and management**
- Look for **potential partners through its contacts**
- VECs of **friends of friends**

VEC formation as a coalitional game

- **Consumers** are **selfish**
- **Coalitional game**:
  - Agents join a coalition if this is in their best interest

Challenges to address

- How do we evaluate a VEC
- How do we build feasible coalitions
- How do we form optimal and stable coalitions

Coalitional value metric

- Given an energy coalition:
  - computes the **total estimated payment**
  - **optimizes** the **buying strategy among energy markets**

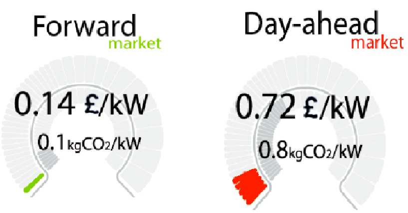

Coalitional value metric

Solves a linear program for a coalition S:

Minimize

$$v(S) = \sum_{t=1}^{N} q_D^t(S) \cdot p_D + N \cdot q_F(S) \cdot p_F$$

t-slot day-head quantity; forward quantity; Day-ahead market price; Forward market price

Subject to

$$q_D^t(S) + q_F(S) \geq e_S^t \quad \forall t = 1 \ldots N$$

Expected demand at slot time t

Coalitional value metric

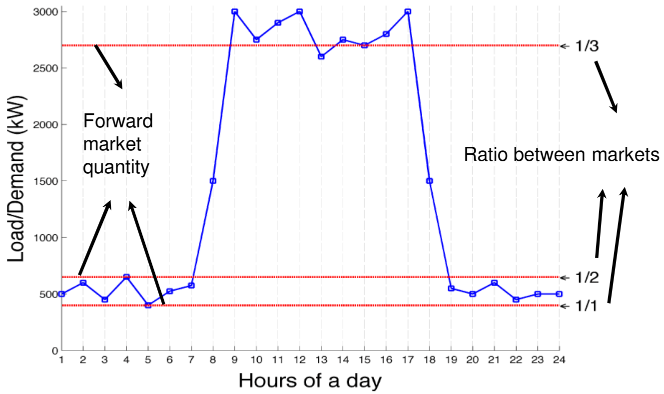

Energy coalitions can buy a continuous amount even when they are not expected to use it all hours of the day

Challenges to address

- How do we evaluate a VEC
- How do we build feasible coalitions
- How do we form optimal and stable coalitions

Coalition enumeration

- **Feasible coalitions** are **restricted by** the **social network graph**
- **Enumerate** all connected sub-graphs [Gutin et al 2008]

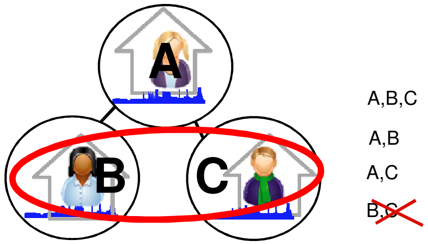

Challenges to address

- How do we evaluate a VEC
- How do we build feasible coalitions
- How do we form optimal and stable coalitions

Coalition Structure Generation

- Aim: identify the set of non-overlapping coalitions with maximal value
- NP-Hard
- Binary integer problem formulation (IP)

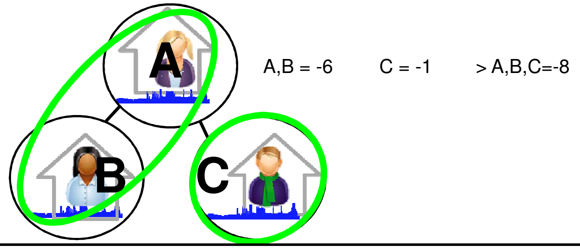

Core-Stable Payoff Distribution

- Find core-stable payments
  - agents have no economical incentive to deviate from optimal coalitions
- Given optimal coalitions use LP formulation

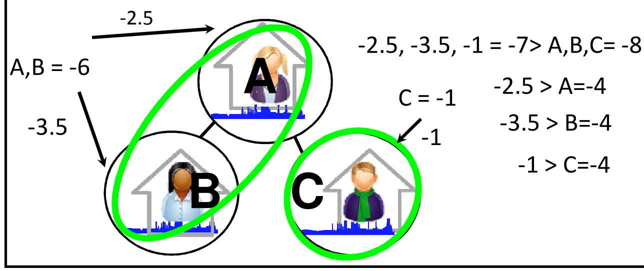

Empirical evaluation

- Real energy profiles from houses in UK
  - Energy consumption averaged over a month
- 20 agents
- Analyze average user gain and coalition structure:
  - network structure (Random, Scalefree and Small-World)
  - # friends acquaintances (#edges/#nodes)
  - Different market conditions

Market (€/KWh)	M1	M2	M3
p_F	70	60	40
p_S	80	80	80

User gain and stability

Topology	Density	% Average Gain			% Empty Core		
		M1	M2	M3	M1	M2	M3
Random	Low	0.5	1.1	6.6	0	2	0
	Medium	0.5	1.3	7.2	48	31	58
	High	0.6	1.3	7.1	54	48	59
ScaleFree	Low	0.6	1.4	7.2	0	0	0
	Medium	0.6	1.4	7	50	40	52
	High	0.6	1.4	7.2	64	60	52

- Lower forward-market price => higher gain
- Higher network density => slightly higher gain, many unstable coalitions
- Similar considerations for small-world

Structure of coalitions

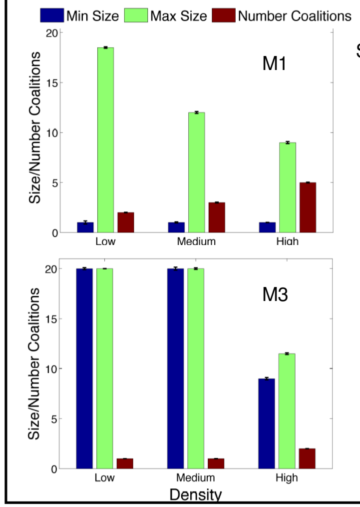

Scale Free

- In M1,M2 coalitions of middle size for all configurations
- In M3 **much larger coalitions**
- **Coalition structure very sensitive to market prices**

Conclusions

Intelligent Agents and MAS:

- "the network is the computer"
- Highly interdisciplinary fields
- Strong focus on building systems
- Many exciting applications