CM30174 + CM50206
Intelligent Agents

Marina De Vos, Julian Padget
East building: x5053, x6971

Agent-Based Modelling / version 0.4

November 29, 2011

Why do ABM?

- Recall institutions: empirical evaluation of institution design
- *In silico* is cheaper than *in vivo*
- Good for feasibility studies: technology, policy, governance
- Get statistics to do the work: scale $\rightsquigarrow$ observation of trends
- Visual interpretation: hides/reveals behaviour

Content

1. Agents or Equations?
2. Case studies
   - School selection
   - Carbon Footprint
   - Call routing
   - Wireless Grids
   - Autonomous vehicles
3. Tools

Objectives

- Illustrate the range of application of agent-based simulation
- Identify problems arising from the approach
- Contrast ABM and equational modelling
- Demonstrate how institutions combine analytical and empirical approaches
- Demonstrate the need for informative visualizations to interpret collective behaviour

Content

1. Agents or Equations?
2. Case studies
3. Tools

Why agent-based simulation?

- We can design mechanisms and institutions
- We can verify institutions — analysts!
- But how do we test them? — empiricists!
- Simulation allows us to evaluate the designs empirically
- But it is not without risk: we have to model precisely enough for the results to be valid
- Agent-based modeling is a bottom-up approach using on local interaction.
- Allows study of mechanics of
  - micro-macro relationships in model and
  - trajectories taken to reach equilibria

How can ABM help?

- Modelling and validating normative frameworks
  - ... or social institutions
  - ... or governance mechanisms
- Populations can take many forms:
  - ... equational
  - ... agent-based (interaction rules, e.g. Life?[1])
  - ... AI-agents (logic, planning, reasoning)
- Institutions too:
  - ... explicit: regulatory or regimented specifications
  - ... implicit: observable through agent (inter-)actions

[1] http://en.wikipedia.org/wiki/Conway's_Game_of_Life

Agent-based simulation

- Comprises agents + environment
- Agents have states and behavioural rules
- Fixed states are parameters and dynamic ones are variables
- Environment may be spatial (e.g., a rectangular grid), or non-spatial (e.g., an abstract trading community)
- Interactions can be direct, where an action immediately changes the state of a partner, or indirect, where an action changes the environment, which, in turn, causes a partner's state to change.
- Environment may be active, having own behaviour to model co-evolution with agents, or passive

Cost of ABM

- Bottom-up $\Rightarrow$ behavioral rules for each agent
- Computational cost higher than calculating dynamics of aggregate global variables of equational models.
- ABMs typically do not contain pro-active, AI-type agents, because:
  - Consumes significant computational resources
  - Full agency makes the system harder to understand — conflicts with aim of scientific experimentation
  - The inherent multi-threaded nature of AI-agency inhibits replication of results — a basic requirement for scientific research.
  - But sometimes need that complication

Is the simulation right?

- Action depends on purpose:

  validation (of hypotheses) vs. prediction

- Four complementary approaches:
  1. Docking: process of aligning the outputs of one simulation with another for given scenarios
  2. Parameter sweep: process of varying a parameter over a range and collecting and visualizing the data to determine the influence of a given paramter
  3. Hypothesis formation and testing: running the simulation to provide evidence for or against hypothesis
  4. Validation against empirical data: are the model outputs sufficiently similar to real-world observations?

Equations vs. Agents 1/2

- Equations model relationships between observables: encoded in the model inputs
- Agents model individual behaviour: relationships emerge as model outputs
- ’What-if’ experiments by changing agent behaviour
- Equations model system-level observables
- Agents model individual observables
- Equations typically regard population as homogeneous
- Agents model indivduals each with potentially different behaviours

Equations vs. Agents 2/2

- Is variation not averaged out in a large enough population?
  Yes, but lose capability to observe individual agent behaviour
- Agents can model more complex situations than equations:
  adding another agent or another attribute is simple
- Extending an equation decreases analytic tractability
- Equations permit proof of mathematical properties
- Agents generate data that constitutes evidence for/against a hypothesis

Summarized from [Parunak et al., 1998]

Agents or Equations?

Ab initio:

- What do you want to model? *big picture or individual interactions?*
- What can you model? *macro or micro relationships?*
- What do you understand? *what behaviour is ($\approx$)certain?*
- What data is available to support/deny hypotheses? *can relevant indicators be collected?*

But, if a model exists, so much the better!

- use it to validate new model
- use new model to validate it

Answer: Agents **and** equations

Content

1 Agents or Equations?

2 Case studies
- School selection
- Carbon Footprint
- Call routing
- Wireless Grids
- Autonomous vehicles

3 Tools

Case studies

1. Social policy analysis: the Baker school reforms (UK, mid 1980s)
2. Evolution of the carbon footprint of the UK housing stock
3. Call routing in call centres
4. Wireless grids

Systems Dynamics

- Systems Dynamics (SD) is widely used in studying complex systems
- SD models identify system variables and describe their dynamics as flows
- Flows take the form of high-level aggregate equations, usually ordinary or partial differential equations, hence equation-based modelling or EBM
- SD model is a set of equations, and execution consists of evaluating them. Good for centralized models of homogeneous entities
- whereas ABM suits domains with a high degree of heterogeneity, localization and distribution.

Quantitative System Dynamics

- Tool for the analysis of dynamic inter-dependencies
- Methodology:
  1. Map processes and lines of influence
  2. Label positive (re-enforcing) or negative (dampening)
  3. Identify sub-systems within the map where all the lines are positive — explosive growth
  4. Likewise negative — implosive collapse
  5. Known as “runaway loops”
- Three questions:
  1. How positive is positive? How fast will system runaway?
  2. How well connected is the sub-system to the driver variables? *Determines system sensitivity to runaway loops*
  3. What opportunities are there to dampen the runaway loops?

QSD Model of UK School Policy

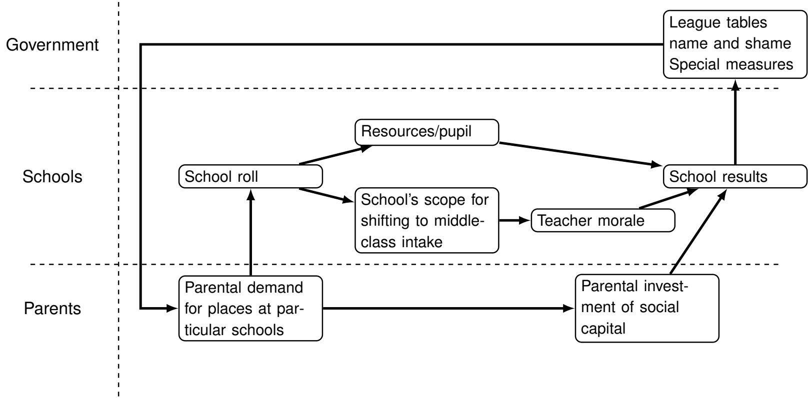

Adapted from [Room and Britton, 2006]

3 class-sensitive schools

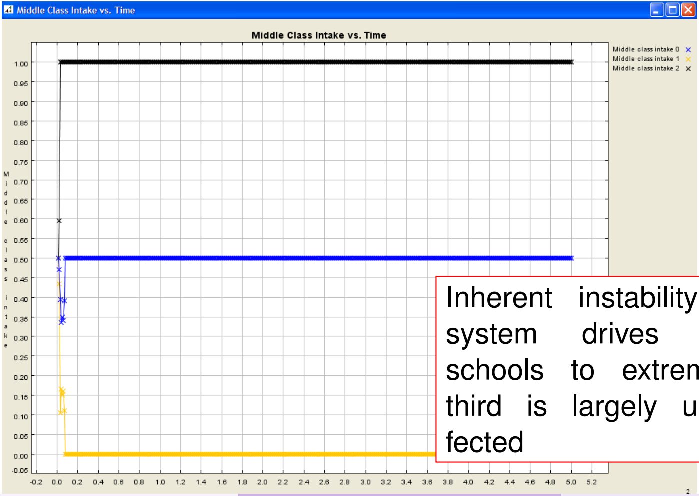

Inherent instability of system drives two schools to extremes, third is largely unaffected

10 class-blind schools

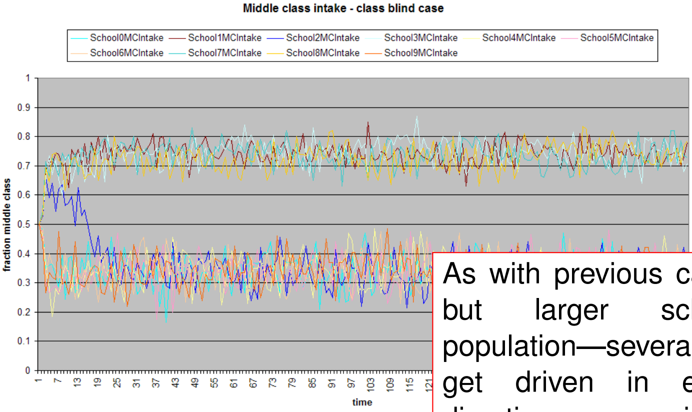

As with previous case, but larger school population—several get driven in each direction, no middle ground

Stochastic shock succeeds

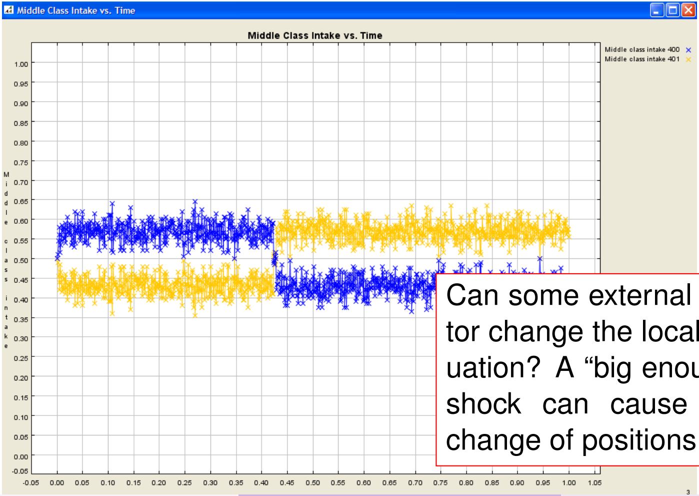

Can some external factor change the local situation? A “big enough” shock can cause exchange of positions

Class blind niche

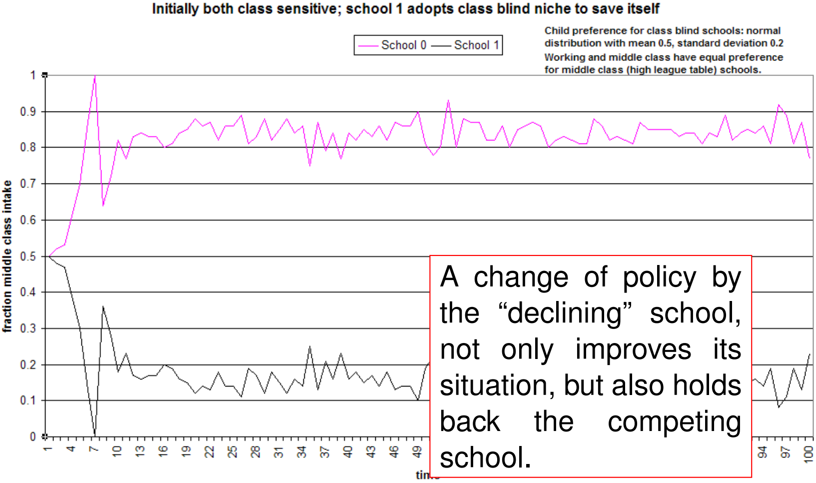

A change of policy by the “declining” school, not only improves its situation, but also holds back the competing school.

Implementation

- Repast
- Agent behaviour expressed as rules using JBOSS rules — standard RETE expert system shell in Java

Reflections on school choice model

- EBM helped validate ABM
- ABM identified assumption in Room-Britton model
- Stepping outside two-school scenario reveals unexpected results: emergent properties or modelling errors?
- ABM permits scenarios that are impossible to analyse in EBM: again are results reliable?

Acknowledgements: Perdita Robinson (CS, 2007), Graham Room (Centre for Social Policy Research)

Exercise: Intelligent vehicles

- **Groups:** 2–3 people
- **Objective:** Sketch a simulation scenario for autonomous vehicles to use ad-hoc networks to organize themselves
- **Plan:**
  - Pair up
  - Core activity [10 mins in all]
    - Identify potential scenarios
    - Choose one to explore in more detail
    - Consider what information is needed (sources) and what communication is required
    - Identify expected outcomes
    - Repeat as desired
  - Reflect and discuss [10 mins]

Carbon Footprint Evolution

“In 2004, more than a quarter of the UKs carbon dioxide emissions a major cause of climate change came from the energy we use to heat, light and run our homes. So its vital to ensure that homes are built in a way that minimises the use of energy and reduces these harmful emissions.” (Communities and Local Government, 2008)

- Use ABM to explore the environmental impact of changes to the UK housing stock
- DECarb [Natarajan and Levermore, 2007]: EBM of transformation of housing stock
- Validation by *back-casting*: like fore-casting, but backwards! From 1996 to 1970.
    - Within 0.9% of actual carbon emissions
    - Within 5.4% of actual energy consumption

Validation + Extension

- Objectives:
  - ABM of housing stock using DECarb front-end
  - Validation by back-casting
  - NEW: Detailed demolition model
  - NEW: Energy-related behaviours
  - NEW: Influence of government policy

DECarb

The user can define the scenario they wish to explore using a series of malleable graphs

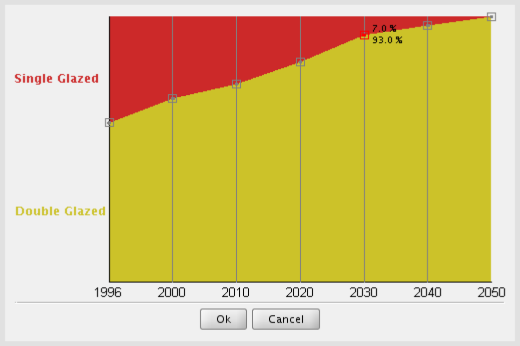

Modelling the UK Housing Stock

- Every household in the UK can be modelled as an individual entity—an agent
- Due to computational resources, every agent currently represents around 200 households
- Potential to model every household with individual behavioural characteristics
- Marionettes: ABM technique, where behaviour is defined globally, but each agent has local state

Results Obtained Using Marionettes

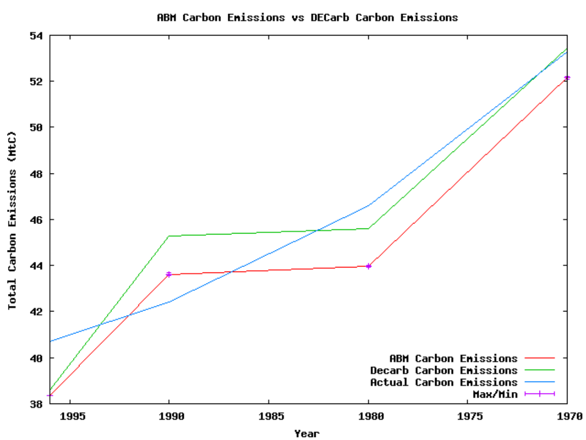

Results Obtained Using Marionettes

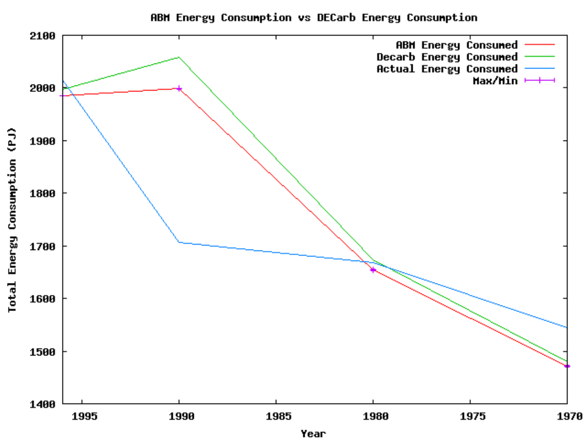

Reflections on carbon footprint model

- EBM helped validate ABM
- ABM also helped identify some anomalies in EBM
- ABM permits exploration of scenarios that are infeasible to model using existing DECarb model
- ABM permits modelling heterogeneous populations of behaviours with the capacity even for individual variation

Acknowledgements: Liam Elliott (CS, 2008), Sukumar Natarajan (Architecture). More details in [Natarajan et al., 2011].

Call routing in call centres

- Fundamental to the operation of most large organisations
- And also emergency services and government agencies
- Function: route calls, monitor KPIs and collect data.
- Aim:
  - Forecast future call volumes
  - Allocate shifts efficiently
  - Experiment with business models
  - Optimize performance + Maintain cost/service tradeoff
- Challenges: poor QoS, high staff turnover, arising from
  - Long waiting queues
  - Inexperienced operators
  - Inaccurate call allocations
  - Inefficient management of staffing levels

Conventional architecture

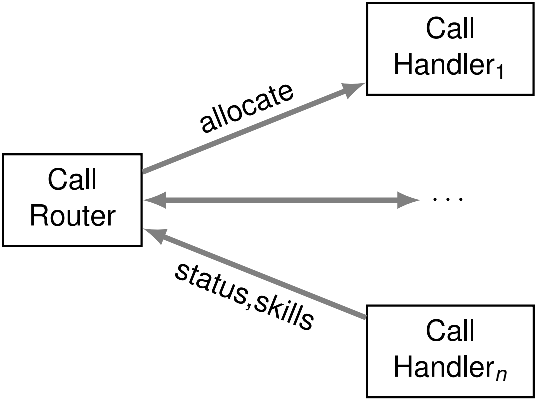

Perception and Reality

- Human view: but modelling directly $\rightsquigarrow$
  - Complex protocols
  - Large state spaces
  - Hard-to-maintain agents
  - Complex call router
  - Centralized decision-making, loss of resilience
- Agent view: individuals that
  - Play roles
  - Function as a collective
  - Distribute work among themselves
  - Implement *observably* the organization

Hierarchical model (ICD)

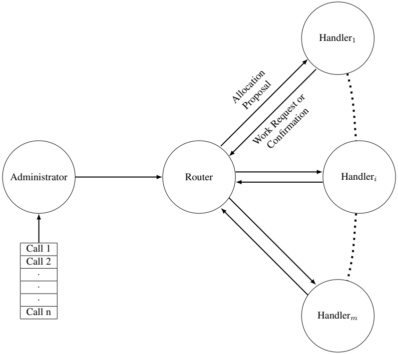

- JADE: complex FSMs, not scalable, not robust
- Cougaar: 560 call handlers processing 43,365 calls over a (simulated) day
- Docks with Call Centre Workshop (CCW) simulator, but (much) slower
- Synthetic and empirical data (Sun Alliance, HSBC, LLoyds, Virgin Mobile)

Self-organizing model (IRN)

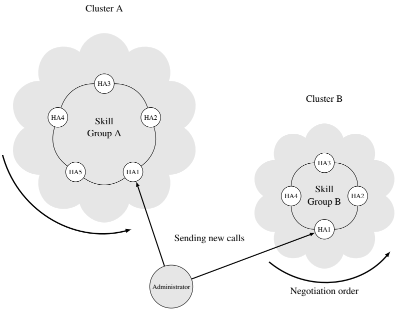

1. Administrator sends call to skill group
2. Skill group identifies handler
3. Or queues call for next available

Simple, inefficient, non-resilient... but satisifies KPIs!

Key Performance Indicators (synthetic data)

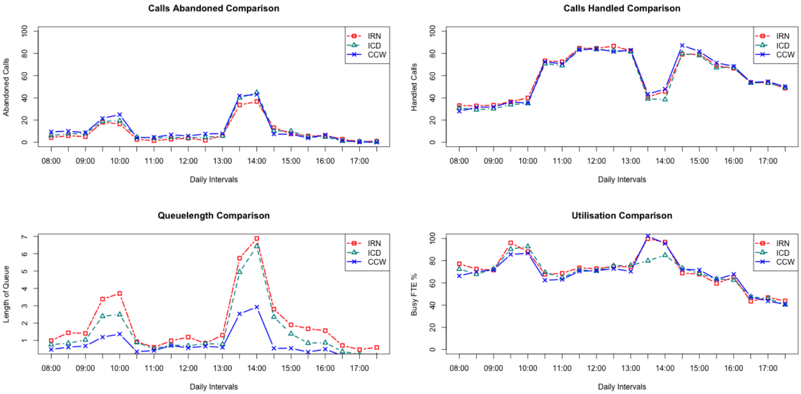

Overall: agent models appear to perform similarly and track CCW

Service levels (synthetic)

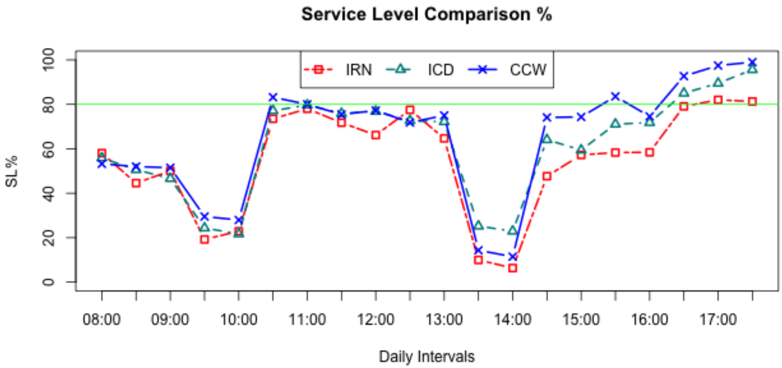

Basket metric that combines previous four

Service levels (actual)

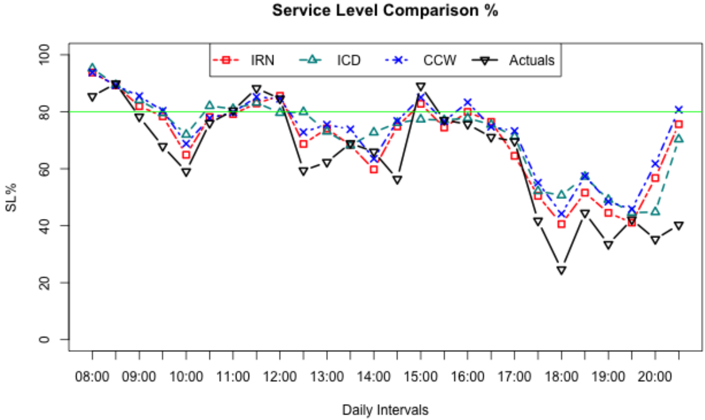

Reflections on call routing

- ABM shows self-organization is a viable alternative: within 5% of CCW on service level
- Too easy to make agents too complicated $\rightsquigarrow$ system lock-up
- Direct modelling of human organizations does not always make the best use of software agents
- Better to build equivalent models than facsimilies?
- Potential to simulate and control with the same sytem

Acknowledgements: Dimitris Traskas (CACI Ltd.). More details in [Traskas and Padget, 2011].

Scenario

- Next generation mobile phones (4G)
- Problem: higher demands, same infrastructure
- Solution? use handsets as part of network
- Benefits:
  - Faster download times: split content, downloading subset with 3G, get rest with wifi from neighbouring handsets
  - Extend battery cycle: trade off high-cost 3G for low-cost wifi communication
  - Reduced load on infrastructure network
- Test case: digital content to distribute to a several nodes that also have a cheap (in terms of power and money) connection via an ad-hoc network

Off-line model

- Focus on static properties of normative system (useful for verification and design of protocols)
- Fast to build, but high chance of over-specification of constraints
- Assumption of limited autonomy of actors
- Starting point for on-line model
- Initial problem:

```
Handset: alice bob
Chunk: x1 x2 x3 x4
Channel: c1 c2
Time: 1 2 3 4
```

- Off-line specification $>$ 150 lines

Visualization

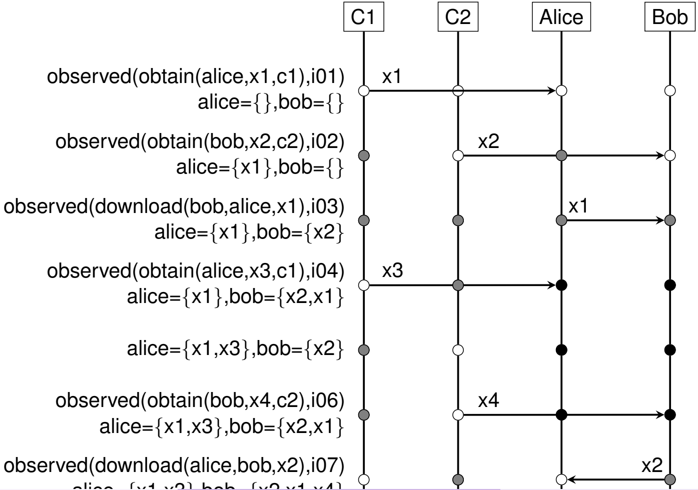

On-line model

- Focus on assisting the running of and adherence to a protocol
- Inclusion of autonomous participant that can reflect upon a normative state
- More realistic with regard to open systems
- More complex and harder to build
- ASP queries take time, but provide essential information:
  - about current state, including applicable norms
  - potential impact of own actions
  - what might happen in the future

On-line sharing specification

```
download(A,X,C) generates intDownload(A,X,C);

intDownload(A,X,C) initiates hasChunk(A,X);
intDownload(A,X,C) terminates downloadChunk(A,X);
intDownload(A,X,C) terminates perm(download(A,X,C1));

send(A,X) generates intSend(A) if hasChunk(A,X);

intSend(B) initiates perm(intReceive(B,X));

send(A,X) generates intReceive(B,X);

intReceive(A,X) initiates hasChunk(A,X);
intReceive(A,X) terminates perm(intReceive(A,X));
intReceive(A,X) terminates pow(intReceive(A,X));
```

The Online Reasoning Process

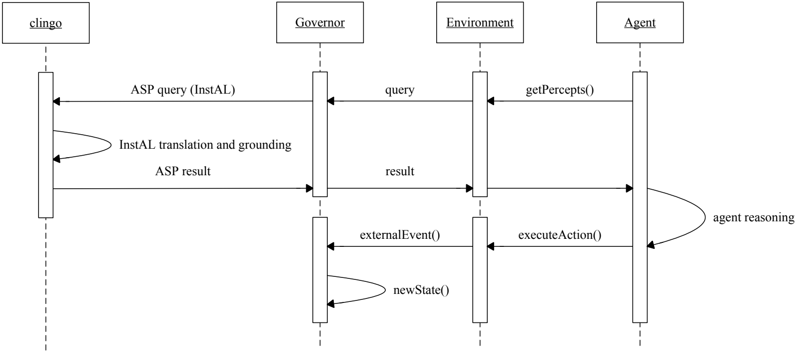

Figure: Interaction of the components

Reflections on Wireless Grids

- AI-type agents
- Use of Jason agent platform (Agentspeak)
- Awkward connection to to institutional model (ASP, clingo)
- Agent behaviour can be affected by institution $\rightsquigarrow$ “what-if” policy experiments

Acknowledgements: Tina Balke (Uni. Bayreuth). More details in [Balke et al., 2011].

Autonomous vehicles

- Objectives:
  - Situational awareness for agents
    - What do you sensors tell you?
    - What do other agents tell you?
  - To establish collective behaviours
  - To work out how much information to reveal
  - To experiment with institutional models in a dynamic environment

Implementation + visualization

- Jason platform – BDI agents
- Tankcoders – networked 3D virtual environment
- Convoy formation:
  - Obstacle detection
  - 2-car convoy
  - 5-car convoy
- Replace simulated cars by Lego robots

Content

1 Agents or Equations?

2 Case studies

3 Tools

Repast

http://repast.sourceforge.net/

- Repast (REcursive Porous Agent Simulation Toolkit)
- Offers a relatively simple Java API for the construction and monitoring of discrete-even simulations
- Extend the class <name> to make different kinds of agents
- Override the `step` method to define the agent's actions
- Examine the state of other agents by
- At each cycle of the simulation, the `step` method of each agent is called.
- Technology is relatively straightforward: challenge is in defining the right experiments and drawing appropriate conclusions.

NetLogo

http://ccl.northwestern.edu/netlogo/

- Written in Java
- Targetted at social science simulations
- Features
  - User programs in a dialect of Logo extended to support agents
  - Can link agents to make aggregates, networks, and graps
  - Cross-platform reproduciblity
  - Visualization of environment in 2D and 3D, interface builder
  - Speed control
  - Extensive model library

Mason

http://www.cs.gmu.edu/~eclab/projects/mason/

- **M**ulti-**A**gent **S**imulator **Of N**eighborhoods
- Claims to be a fast discrete-event multiagent simulation library core in Java
- Extensive model library
- Visualization in 2D and 3D
- Support for checkpointing and migration
- Reproducibility across platforms

Summary 1/2

- Why use ABM?
  - Allows modeller to concentrate on interactions between components: bottom-up
  - Ease of modification/extension: new behaviour, additional events
  - Heterogenous populations
- Why not to use ABM!
  - Results are empirical not analytical: evidence not proof
  - Validation is difficult
  - Loss of perspective: need top-down approach too

Summary 2/2

- Attractive method for modelling/exploring mechanism design
- Tradeoff: simple model but lots of run-time — plan experiments carefully
- Possibility of exploring mixed human/simulation environments using avatars (participatory simulation)
- But easy to generate unsound results — and bugs are hard to spot!

Recommended Reading

- Wooldridge: does not discuss ABM
- [Parunak et al., 1998] compares equational and agent based simulation
- [Gilbert and Bankes, 2002] gives a brief survey and evaluation of software platforms for ABM
- `www.pnas.org`, May (suppl. 3), 2002 has a collection of papers about agent-based modelling

References

Balke, T., Vos, M. D., and Padget, J. (2011).
Analysing energy-incentivized cooperation in next generation mobile networks using normative frameworks and an agent-based simulation.
Future Generation Computer Systems, 27(8):1092–1102.

Gilbert, N. and Bankes, S. (2002).
Platforms and methods for agent-based modeling.
Proceedings of the National Academy of Sciences, 99 (suppl. 3):7197–7198.
Available via `www.pnas.orgcgidoi10.1073pnas.072079499`.

Natarajan, S. and Levermore, G. (2007).
Predicting future UK housing stock and carbon emissions.
Energy Policy, 35(11):5719–5727.

Natarajan, S., Padget, J., and Elliott, L. (2011).
Modelling UK domestic energy and carbon emissions: an agent-based approach.
Energy & Buildings, 43:2602–2612.
`http://dx.doi.org/10.1016/j.enbuild.2011.05.013`.

Parunak, H. V. D., Savit, R., and Riolo, R. L. (1998).
Agent-based modeling vs. equation-based modeling: A case study and users' guide.
In Sichman, J. S., Conte, R., and Gilbert, N., editors, *MABS*, volume 1534 of *Lecture Notes in Computer Science*, pages 10–25. Springer.

Room, G. and Britton, N. (2006).
The dynamics of social exclusion.
International Journal of Social Welfare, 15.
DOI:10.1111/j.1468-2397.2006.00417.x.

Traskas, D. and Padget, J. (2011).
A multi-agent systems approach to call-centre management.
International Journal of Parallel, Emergent and Distributed Systems, 26(5):347–367.