Travelling collectives: The human side of transportation modelling

Michael Rovatsos
School of Informatics, University of Edinburgh

SICSA Workshop on modelling and optimisation of real-world transportation
University of Stirling, 16th January 2015

Social Computation

- Large-scale man-machine collectives that perform complex computations mediated by digital media
- Often based on crowdsourcing and/or human computation, e.g. Captcha, Galaxy Zoo, TaskRabbit, Kickstarter
- Offers a new vision for collectively intelligent systems, where people and artificial agents complement each other
- Existing systems have been built in ad hoc fashion, achieving good design rather a "black art"
- Many applications in transportation modelling and optimisation

SEVENTH FRAMEWORK PROGRAMME

The SmartSociety Project

- 4-year €6.8M EU FP7 FET Integrated Project, co-ordinated by University of Trento
- Aim: building hybrid and diversity-aware collective adaptive systems to solve challenging societal problems
- Edinburgh focus within the project: **social orchestration & compositionality** of social computation systems
  - lightweight architectures for composing Web workflows
  - enabling "compositionality" with human and machine help
  - identifying emergent structure and evolving system design

SEVENTH FRAMEWORK PROGRAMME

Focus of this talk

1. Algorithms
2. Architectures
3. People

1. Algorithms

Automated Planning

- Planning is a key ability of intelligent systems
  - Conceptually: how do I achieve a goal given what's true now and what parts of the world I can influence?
  - Algorithmically: the problem of generating action sequences that will reach a state with certain properties
- Relevance for social computation: useful to describe, synthesise, verify, and execute complex tasks
- In the standard "single-agent" case, there exist highly scalable and domain-indepdendent algorithms
- In the "multiagent" case, this get much more complicated

Multi-Objective Planning

- Agents have independent, potentially conflicting goals – *this makes things much harder!*
- Strategic problem, solution concepts based on notions of stability and equilibrium
- Problem depends on whether contracts can be enforced and utility can be transferred
- Like concurrent planning with additional constraints on plan cost to individuals
- Hard to define meaningful solution concepts if goals incompatible or agents untrustworthy

SEVENTH FRAMEWORK PROGRAMME

Example

- Delivery domain

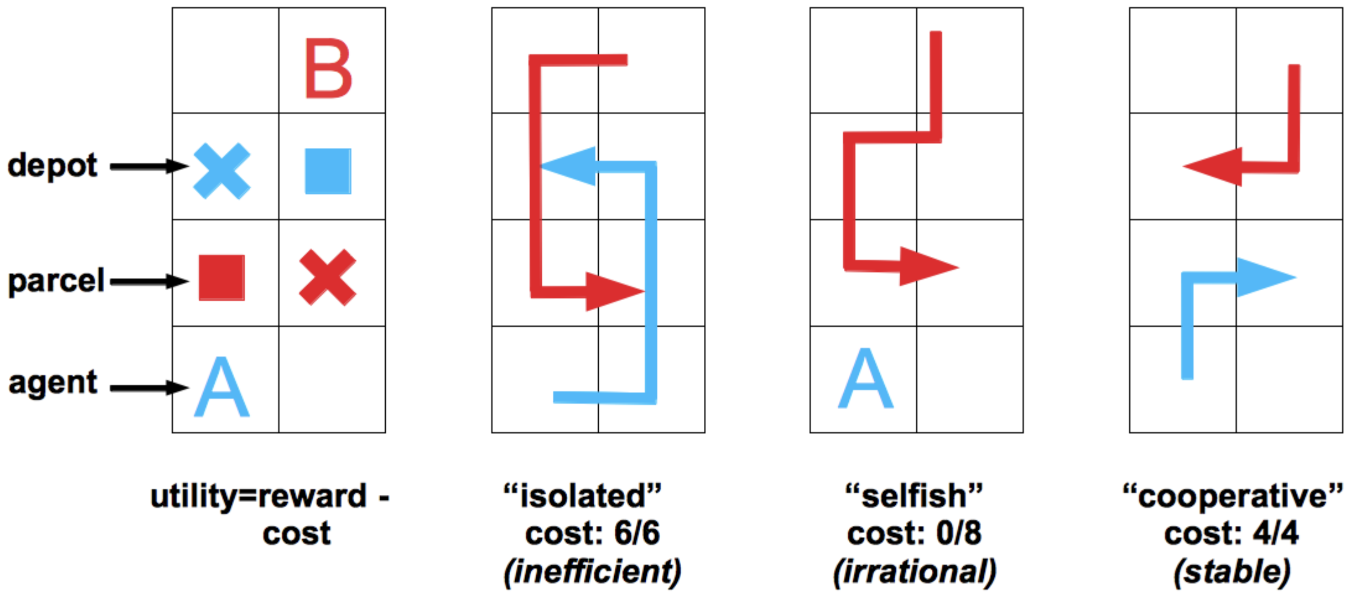

- We focus on domains where enumeration of plans is impossible – i.e. this is not standard game theory

Planning for Self-Interested Agents

- Best-Response Planning (Jonsson & MR):
  - iterative method of optimising agents' individual plans without breaking others' plans
  - computes equilibrium plans fast in congestion games, restricted to interactions regarding cost
  - useful for plan optimisation in unrestricted domains
- Later also extended to include "compress-and-expand" algorithm to produce initial concurrent plan
  - only for domains where agents can achieve their individual goals alone; but useful for plan cost optimisation

SEVENTH FRAMEWORK PROGRAMME

Application: Ridesharing

- Hrncir & Hlavaty's system uses BRP to calculate travel routes using real-world UK public transportation data and private cars (>200,000 connections)

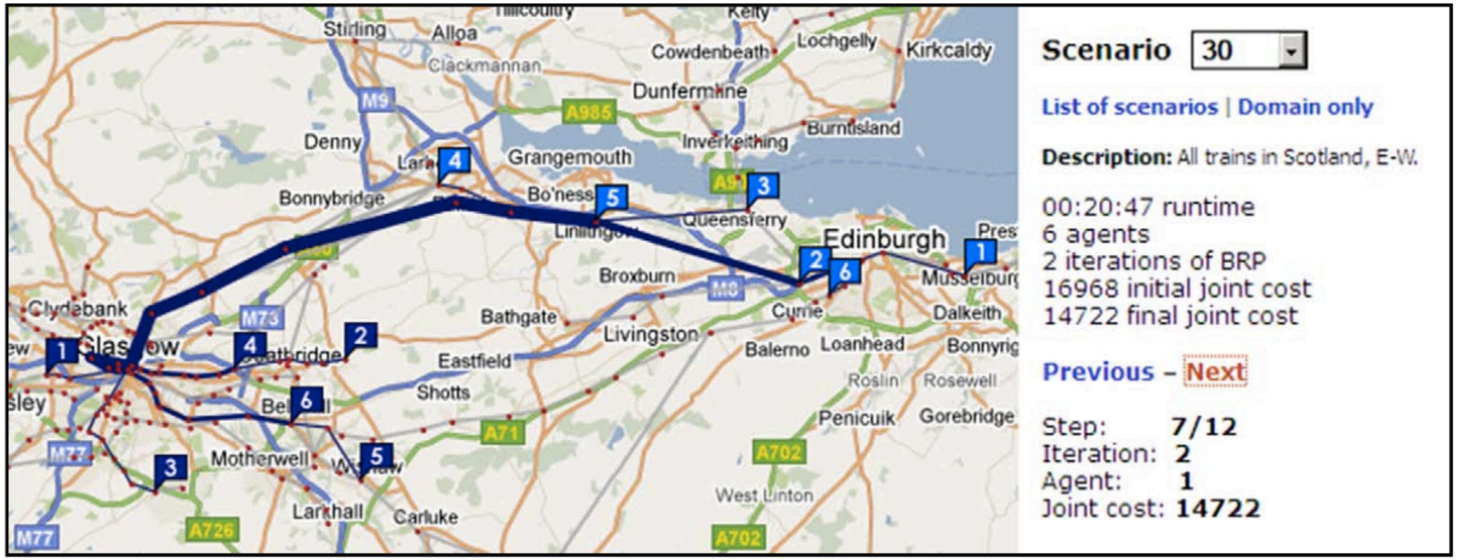

Relevance to transportation

- Compute travel routes that are acceptable for all travellers involved
  - reduces probability of individuals deviating from joint plan (though collusion among sub-groups is still possible)
  - allows for broad range of cost functions and taking account of system-level goal (e.g. to influence traveller behaviour or avoid congestion)
  - may help convince people of benefits of ridesharing by helping to navigate unmanageable space of possible rides
- The same algorithms can be applied to many problems other than ridesharing (e.g. logistics, road planning, etc)

2. Architectures

Deploying algorithms on the Web

- Large-scale agent populations interacting sporadically
- Continually elicited human knowledge and preferences
- Require robust, scalable, and fast implementations

Social orchestration

- Discovery: peer as service provider/task requester
- Composition: creation of possible tasks
  - Tasks may involve several peers and several steps
- Recommendation: presenting the right tasks to peer groups
  - Heuristics for filtering are informed by past data and feedback
  - Incentives added to increase acceptability
- Negotiation: getting agreement for tasks
- Execution (specified, but not implemented) and feedback

Adaptive social orchestration

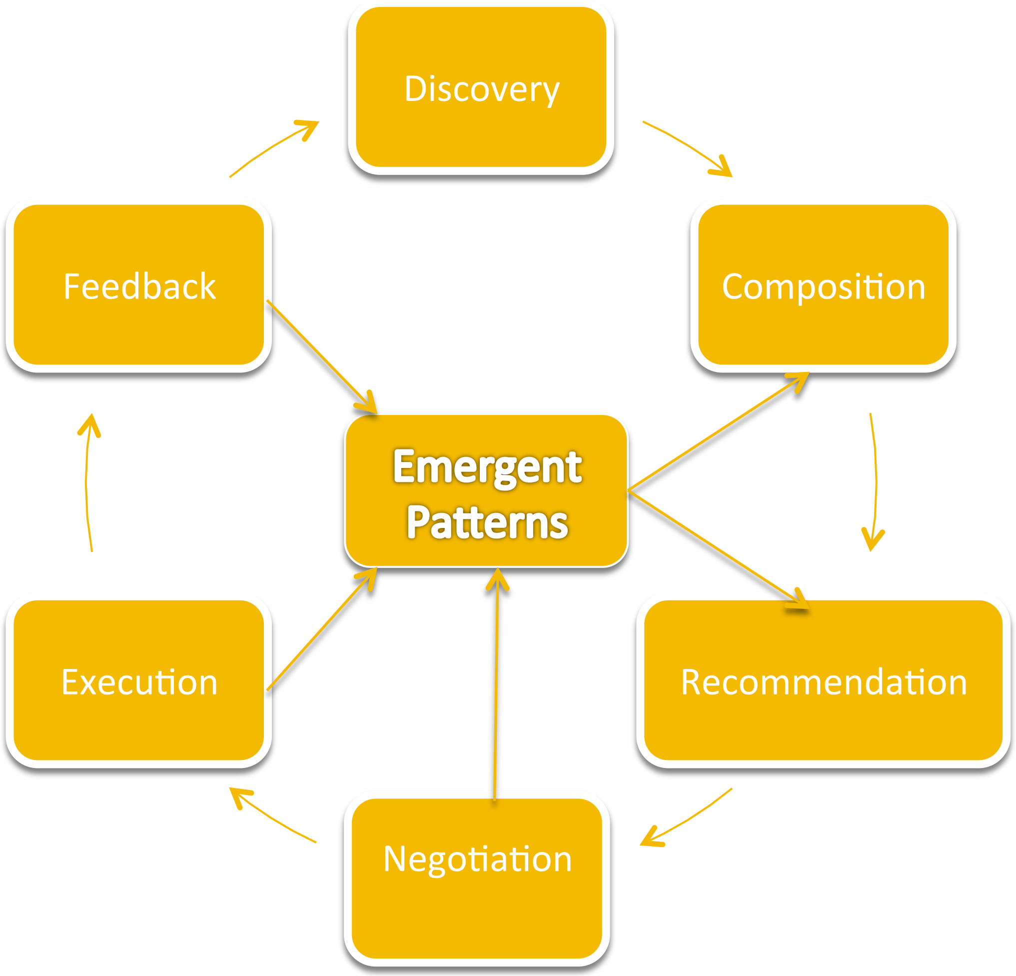

Group task recommendation

- What is different from traditional recommender systems?
  - We don't know whether a solution exists for a requested objective a priori (cannot just propose nearest "product")
  - Impossible to compute all possible solutions offline (and annotate them for retrieval), computation takes (run-)time
  - We require agreement of all parties for a task to happen, i.e. solution must rank high on everyone's preferences
  - Individual preferences highly context-dependent as they depend on concrete team partners proposed by a solution
  - Data obtained from negotiation/execution/feedback refers to whole teams (correlated views), not just individuals

3. People

The human in the loop

Many people don't

- care about optimality, but a lot about safety
- want to be burdened by being given too many options
- have fixed (or even consistent) preferences
- want to negotiate or compromise
- want to delegate decisions to machines

This means that much of what we emphasise on the computational side may not be that relevant!

Conclusions

- Described intelligent transportation work we have done in the context of social computation work
- Algorithmic problems challenging, and closely intertwined with architectural system design considerations
- Human-centric engineering needs to account for many factors often ignored in traditional computational models
- There are many opportunities for engaging with this issue, but they require working with users, not just machines!

SEVENTH FRAMEWORK PROGRAMME

Thank you for your attention!

- Questions?

Abstract

In recent work, we have started to look at transportation planning for large-scale populations of travellers, especially through ridesharing. In this talk, we focus on the impact working with actual users has affected our research in interesting an unexpected ways. We start by presenting work on well-defined (but hard) combinatorial problems involved in planning joint rides for groups of self-interested travellers, and then look at the ways in which the reality of dealing with human users challenges some of our theoretical assumptions, while raising interesting new computational problems.

These affect both the design of optimisation and travel recommendation algorithms, as the development of appropriate distributed systems platforms to deploy intelligent travel planning systems in the real world.