Agent-Based Systems

Michael Rovatsos
mrovatso@inf.ed.ac.uk

Lecture 7 – Methods for Coordination

Where are we?

Last time . . .

- Agent communication
- Speech act theory
- Agent communication languages (KQML/KIF, FIPA-ACL)
- Interaction Protocols
- Ontologies for communication

Today . . .

- Methods for Coordination

Methods for Coordination

- Coordination is the process of *managing inter-dependencies between agents' activities*
- Remember our previous definition

  Coordination is a special case of interaction in which agents are aware how they depend on other agents and attempt to adjust their actions appropriately.

- Actually this only covers agent-based coordination, but there can also be centralised mechanisms
- In contrast to cooperation, coordination is also necessary in non-cooperative systems (unless agents ignore each other)

Coordination within interaction

Coordination in a general typology of interaction:

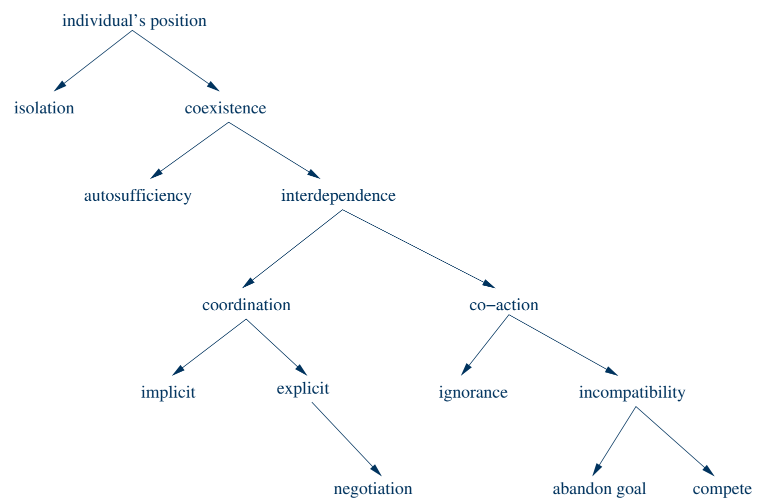

Typology of coordination relationships

- More specific typology in the context of multiagent planning (von Martial, 1990):

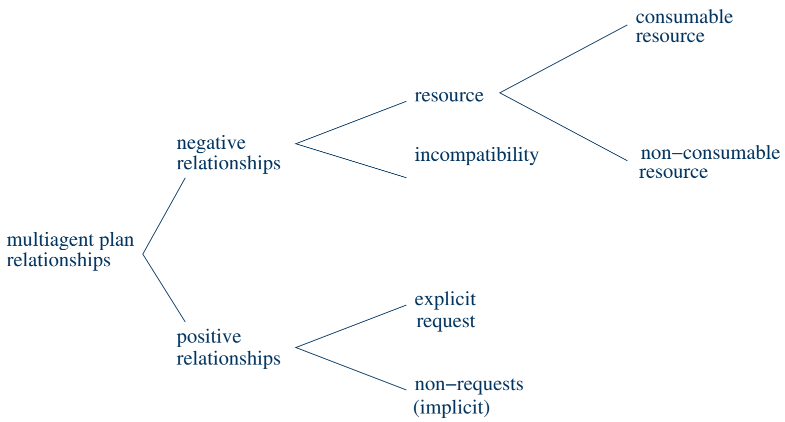

Typology of coordination relationships

- Positive relationships: relationships between two agents' plans for which benefit will be derived for at least one agent if plans are combined
- Requests: explicitly asking for help with own activities
- Non-requested: pareto-like implicit relationships
  - action equality relationships: sufficient if one agent performs action both agents need
  - consequence relationships: side effects of agent's plan achieve other's goals
  - favour relationships: side effects of agent's plan make goal achievement for other agent easier
- Basic difference to traditional computer systems: coordination is achieved at **run time** rather than **design time**
- Remainder of lecture: discussion of different approaches to achieve coordination

Partial global planning

- **Partial global planning** (PGP): exchange information to reach common conclusions about problem-solving process
- Partial – individual agents don't generate plan for entire problem
- Global – agents use information obtained from others to achieve non-local view of problem
- Three iterated stages:
  1. Agents deliberate locally and generate short-term plans for goal achievement
  2. They exchange information to determine where plans and goals interact
  3. Agents alter local plans to better coordinate their activities
- **Meta-level structure** guides the coordination process, dictates information exchange activities

Partial global planning

- Central data structure: partial global plan, containing:
  - Objective: larger goal of the system
  - Activity maps: describe what agents are doing and the results of these activities
  - Solution construction graph: describes how agents should interact and exchange information to achieve larger goal
- Framework extended/refined in **Generalized PGP** (GPGP)
- GPGP introduces five techniques for coordinating activities, i.e. strategies for
  - updating non-local viewpoints (share all/no/some information)
  - communicating results
  - handling simple (action) redundancy
  - handling hard ("negative") coordination relationships (mainly by means of rescheduling)
  - handling soft ("positive") coordination relationships (rescheduling whenever possible, but not "mission critical")

(G)PGP application – DVMT

- Distributed Vehicle Monitoring Testbed (DVMT): one of the earliest testbeds for CDPS networks
- Aim of the system: tracking number of vehicles passing within a range of distributed sensors
- Different problem-solving strategies were successfully tested in this domain using the (G)PGP approach
- Data-driven domain: challenge is to process vehicle movement data to infer their paths in a timely fashion
- Interesting: distributed sensor networks currently a hot topic, this research started in 1980!

Joint intentions

- We discussed intentions in practical (single-agent) reasoning
- But intentions also provide stability and predictability necessary for social interaction
- Therefore also significant for coordination, especially teamwork
- Helps to distinguish between non-cooperative and cooperative coordinated activity
- Basic question: in which way are individual intentions different from (and what role do they play in) **collective intentions**?
- Remember Cohen and Levesque's theory of intentions? They extended it to teamwork situations, introducing a notion of "responsibility"

Joint intentions

- Example: We try to lift a stone together, and I discover it won't work
  ➡ individually rational behaviour: drop the stone
- However, this is not really cooperative (we should at least inform other)
- Two important notions:
  - **commitments** (pledges or promises to underpin an intention)
  - **conventions** (mechanisms for monitoring commitment, mechanics of adopting/abandoning commitments)
- Agents can commit themselves to actions or states of affairs
- Commitments are **persistent**, i.e. they are not dropped unless special circumstances arise
- Conventions define these circumstances, e.g. that motivation for goal is no longer present, that it is or can never be achieved

Joint intentions

- **Joint** commitments have a distributed state among team members
- Conventions describe, e.g. that an agent should inform others when it drops an individual commitment
- Notion of **joint persistent goal** (JPG): A goal φ with motivation (reason) ψ such that:
  - initially all agents don't believe φ but believe it is possible
  - every agent has goal φ until termination condition is satisfied
  - termination condition: mutual belief that φ satisfied, impossible to achieve, or motivation ψ no longer present
- While termination condition is not met, if any agent *i* believes φ is achieved or impossible or that ψ is no longer present it has a persistent goal that this becomes mutual belief until termination condition is met

Teamwork-based model of CDPS

- Practical model of how CDPS can operate using a teamwork approach
- Stage 1: **Recognition** of a goal that can be achieved through cooperation (e.g. an agent can't do it (efficiently) on his own)
- Stage 2: **Team formation**, i.e. assistance solicitation
  - if successful, this results in nominal commitment to collective action
  - deliberation phase, ends in agreement on ends (not on means)
  - rationality plays a role in deciding whether to form a group
- Stage 3: **Plan formation** (joint means-ends reasoning, e.g. through negotiation or argumentation)
- Stage 4: **Team action** with JPG as an example convention that governs joint plan execution

Mutual modelling

- Based on putting ourselves in the place of the other
- Involves modelling others' beliefs, desires, and intentions . . .
- . . . and coordinating own actions depending on resulting predictions
- Explicit communication is not necessary
- MACE one of the first systems to use **acquaintance models** for this purpose
- Acquaintance knowledge involves information about others'
  - **Name** unique to every agent
  - **Class** (group to which agent belongs)
  - **Roles** played by an agent in a class
  - **Skills** as the capabilities of the modelled agent
  - **Goals** that the modelled agent wants to achieve
  - **Plans** describing how modelled agent attempts to achieve goals
- Agent also explicitly models itself!

Norms and social laws

- **Norms** are established patterns of expected behaviour, **social laws** often add some authority to that (can be enforced or not)
- Idea: to strike a balance between autonomy and goals of entire society
- Such conventions make decision making easier for agent
- Can be designed offline or emerge from within the system
- The former is simpler, the latter more flexible
- Hard to predict which norm will be optimal for a system at design time
- But also hard to derive global conventions from agents' point of view given only local information

Emergent social norms and laws

- Example: the t-shirt game
  - agents wear red or blue t-shirt (initially at random), goal is for everyone to wear the same colour
  - agents are randomly paired in each round of the game, get to see other's t-shirt colour, and then may decide to switch colour
- Problem: agent must decide which convention to adopt although no global information is available
- Possible update functions (=decision rules based on history):
  - Simple majority: agent chooses colour observed most often
  - Simple majority with agent types: agents confide in certain other agents and exchange memory with them to inform their decision
  - Simple majority with communication on success: agents will communicate (successful part of) memory if success rate exceeds a threshold
  - Highest cumulative reward: uses strategy that has had the highest cumulative reward so far

Emergent social norms and laws

- All update functions converged to some convention
- Measure: time taken to converge
- Memory restarts were investigated to model "new ideas"
- But also stability important (we don't want society to change conventions all the time)
- Basic result: for highest cumulative result rule, for any $0 \leq \epsilon \leq 1$ agents will reach agreement within n rounds with probability $1 - \epsilon$
- Also, once reached, the convention will be stable
- And convention is efficient, i.e. it guarantees payoff no worse than that obtainable from sticking to initial choice
- Note that change of norm may be expensive in practice!

Offline design

- Closely related to mechanism design
- Formally, remembering our agent model $Ag : \mathcal{R}^E \rightarrow Ac$ we can define **constraints** $\langle E', \alpha \rangle$ where
  - $E' \subseteq E$
  - $\alpha \in Ac$

  such that α is forbidden in any state from E'
- A social law is a set of such constraints, agents/plans are **legal** if they never attempt to perform forbidden actions
- Given a set $F \subseteq E$ of **focal states** (states that should always be allowed), a **"useful social law problem"** is to find a social law that will allow agents to legally visit any state in F
- General problem NP-complete, tractable special cases not realistic

Summary

- Coordination: managing interactions effectively
- Different methods for coordination
- Partial global planning: achieving a global view through information exchange
- Joint intentions: extending the BDI paradigm to include joint intentions, collective commitments and conventions
- Mutual modelling: taking the role of the other to predict their actions
- Norms and social laws: coordination through offline/emergent constraints on agent behaviour
- Next time: **Multiagent Interactions**