

CHAPTER 1: INTRODUCTION

Multiagent Systems

<http://www.csc.liv.ac.uk/~mjw/pubs/imas/>

Five Trends in the History of Computing

- *ubiquity*;
- *interconnection*;
- *intelligence*;
- *delegation*; and
- *human-orientation*.

Ubiquity

- Continual reduction in cost of computing makes it possible to introduce processing power into places and devices that would have once been uneconomic.
- As processing capability spreads, sophistication (and intelligence of a sort) becomes *ubiquitous*.
- What could benefit from having a processor embedded in it?

Interconnection

- Computer systems no longer stand alone, but are networked into large distributed systems.
- Internet an obvious example, but networking is spreading its ever-growing tentacles.
- Since distributed and concurrent systems have become the norm, some researchers are putting forward theoretical models that portray computing as primarily a process of interaction

Intelligence

- The complexity of tasks that we are capable of automating and delegating to computers has grown steadily
- If you don't feel comfortable with this definition of "intelligence", it's probably because you are a human. . .

Delegation

- Computers are doing more for us . . . without our intervention
- We are giving control to computers, even in safety critical tasks
- One example: fly-by-wire aircraft, where the machine's judgment may be trusted more than an experienced pilot
- Next on the agenda: fly-by-wire cars, intelligent braking systems, cruise control that maintains distance from car in front. . .

Human Orientation

- The movement away from machine-oriented views of programming toward concepts and metaphors that more closely reflect the way we ourselves understand the world
- Programmers (and users!) relate to the machine differently
- Programmers conceptualize and implement software in terms of ever higher-level – more *human-oriented* – abstractions

Programming progression

- Programming has progressed through:
 - machine code;
 - assembly language;
 - machine-independent programming languages;
 - sub-routines;
 - procedures & functions;
 - abstract data types;
 - objects;to
 - agents.

Other Trends in Computer Science

- the Grid;
- ubiquitous computing;
- semantic web.

The Grid

- The *Grid* aims to develop massive-scale open distributed systems, capable of being able to effectively and automatically deploy and redeploy computational (and other) resources to solve *large* computational problems:
 - huge datasets;
 - huge processing requirements.
- Current Grid research focussed mainly on *middleware*

The Grid and Multiagent Systems

‘The Grid and agent communities are both pursuing the development of such open distributed systems, albeit from different perspectives. The Grid community has historically focussed on [...] “brawn”: interoperable infrastructure and tools for secure and reliable resource sharing within dynamic and geographically distributed virtual organisations (VOs), and applications of the same to various resource federation scenarios. In contrast, those working on agents have focussed on “brains”, i.e., on the development of concepts, methodologies and algorithms for autonomous problem solvers that can act flexibly in uncertain and dynamic environments in order to achieve their objectives.’ (Foster et al, 2004)

Ubiquitous Computing

‘[P]opulations of computing entities – hardware and software - will become an effective part of our environment, performing tasks that support our broad purposes without our continual direction, thus allowing us to be largely unaware of them. The vision arises because the technology begins to lie within our grasp. This tangle of concerns, about future systems of which we have only hazy ideas, will define a new character for computer science over the next half-century.’ (Milner, 2006)

The Semantic Web

- The semantic web aims to annotate web sites with *semantic markup*: information in a form processable by computer, typically relating to the content of the web site.
- The idea is that this markup will enable browsers (etc) provide richer, more meaningful services to users.

Berners-Lee on the Semantic Web

‘I have a dream for the Web [in which computers] become capable of analysing all the data on the Web – the content, links, and transactions between people and computers. A ‘Semantic Web’, which should make this possible, has yet to emerge, but when it does, the day-to-day mechanisms of trade, bureaucracy and our daily lives will be handled by machines talking to machines. The ‘intelligent agents’ people have touted for ages will finally materialise.’ (Berners-Lee, 1999)

Agents: A First Definition

An *agent* is a computer system that is capable of *independent* (*autonomous*) action on behalf of its user or owner (figuring out what needs to be done to satisfy design objectives, rather than constantly being told).

Multiagent Systems: A First Definition

A *multiagent system* is one that consists of a number of agents, which *interact* with one-another. In the most general case, agents will be acting on behalf of users with different goals and motivations. To successfully interact, they will require the ability to *cooperate*, *coordinate*, and *negotiate* with each other, much as people do.

A Vision: Autonomous Space Probes

When a space probe makes its long flight from Earth to the outer planets, a ground crew is usually required to continually track its progress, and decide how to deal with unexpected eventualities. This is costly and, if decisions are required *quickly*, it is simply not practicable. For these reasons, organisations like NASA are seriously investigating the possibility of making probes more autonomous — giving them richer decision making capabilities and responsibilities.

This is not fiction: NASA's DS1 is doing it now!

A Vision: Internet Agents

Searching the Internet for the answer to a specific query can be a long and tedious process. So, why not allow a computer program — an agent — do searches for us? The agent would typically be given a query that would require synthesising pieces of information from various different Internet information sources. Failure would occur when a particular resource was unavailable, (perhaps due to network failure), or where results could not be obtained.

The Micro and Macro Problems

- *Agent design*: How do we build agents that are capable of independent, autonomous action in order to successfully carry out the tasks that we delegate to them?
- *Society Design*: How do we build agents that are capable of interacting (cooperating, coordinating, negotiating) with other agents in order to successfully carry out the tasks that we delegate to them, particularly when the other agents cannot be assumed to share the same interests/goals?

Some Views of the Field

- *Agents as a paradigm for software engineering*:
Software engineers have derived a progressively better understanding of the characteristics of complexity in software. It is now widely recognised that *interaction* is probably the most important single characteristic of complex software.

- *Agents as a tool for understanding human societies:*

Multiagent systems provide a novel new tool for simulating societies, which may help shed some light on various kinds of social processes.

Objections to MAS

- Isn't it all just Distributed/Concurrent Systems?
- Isn't it all just AI?
- Isn't it all just Economics/Game Theory?
- Isn't it all just Social Science?