LECTURE 1: INTRODUCTION

An Introduction to Multiagent Systems

CIS 7412, Fall 2011

Five Trends

- Five ongoing trends have marked 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
- Many of these tasks are ones that can be thought of as requiring a good deal of intelligence.
- If you don't feel comfortable with this definition of "intelligence", it's probably because you are a human. . .
- Imagine whether you would be comfortable with your pet dog or cat doing these tasks.

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. . .

Oh, wait . . .

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.

- Just as moving from machine code to higher level languages brings an efficiency gain, so does moving from objects to agents.
- A 2006 paper:

  S. Benfield, Making a Strong Business Case for Multiagent Technology, Invited Talk at AAMAS 2006.

  claims that developing complex applications using agent-based methods leads to an average saving of 350% in development time (and up to 500% over the use of Java).

Global Computing

- What techniques might be needed to deal with systems composed of 10^{10} processors?
- Don't be deterred by this seeming to be "science fiction"

  (Hundreds of millions of people connected by email once seemed to be science fiction. . .)
- Current software development models can't handle this. . . Where does it bring us?

Global Computing 2

- Delegation and Intelligence imply the need to build computer systems that can act effectively on our behalf

  This implies the ability of computer systems. . .

  – to act independently

  – to act in a way that represents our best interests while interacting with other humans or systems.

Interconnection and Distribution

- Now core motifs in Computer Science
- But Interconnection and Distribution, coupled with the need for systems to represent our best interests, implies systems that can *cooperate* and *reach agreements* (or even *compete*) with other systems that have different interests (much as we do with other people)
- So Computer Science expands: These issues were not studied in Computer Science until recently All of these trends have led to the emergence of a new field in Computer Science: *multiagent systems*.

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.

Spacecraft Control

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, NASA would like probes that are more autonomous and have richer decision making capabilities and responsibilities.

Spacecraft Control (2)

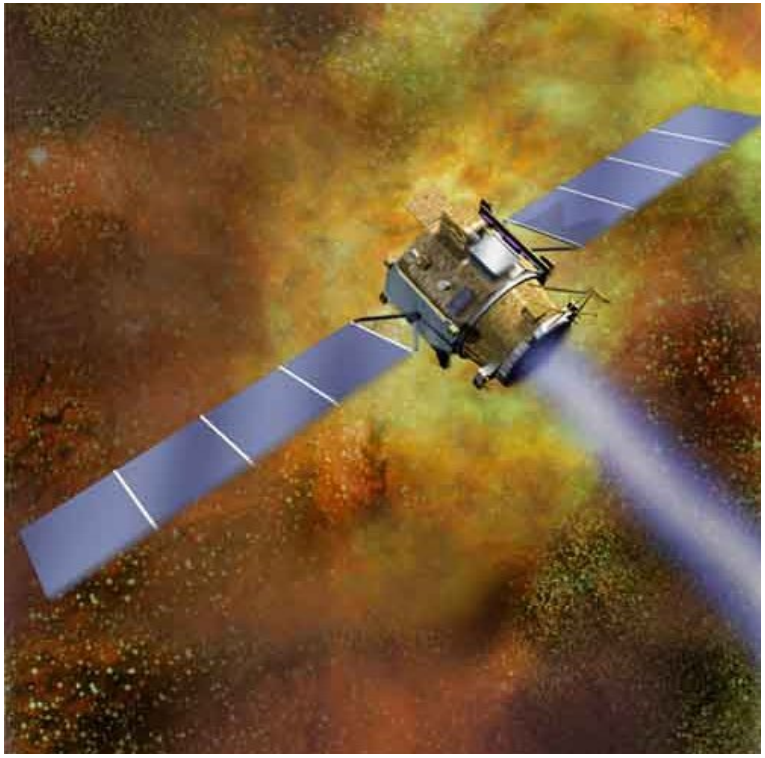

Spacecraft Control (3)

This is not fiction: NASA's DS1 has already shown this is possible!

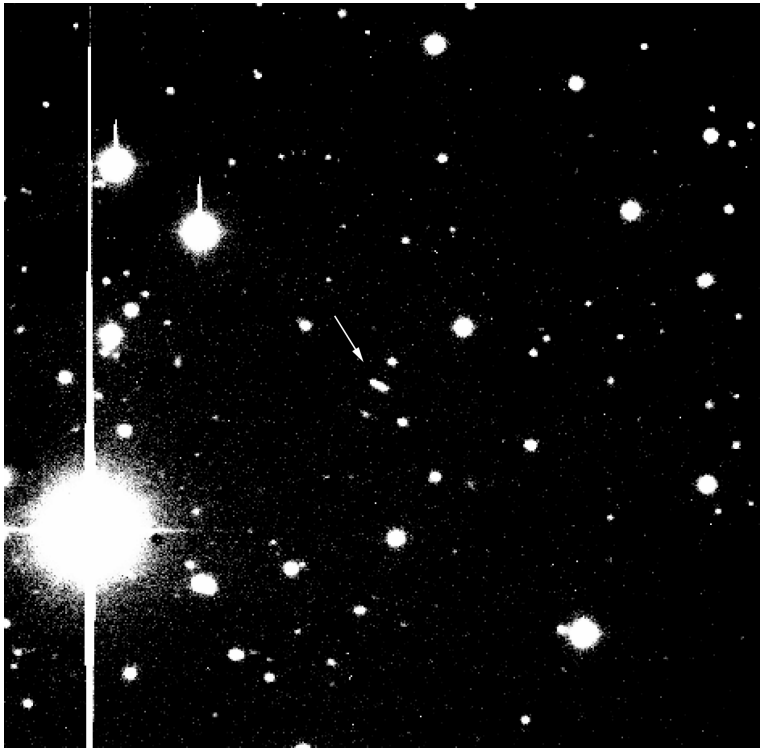

(DS1 seen 2.3 million miles from Earth)

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.

Natural-born cyborgs?

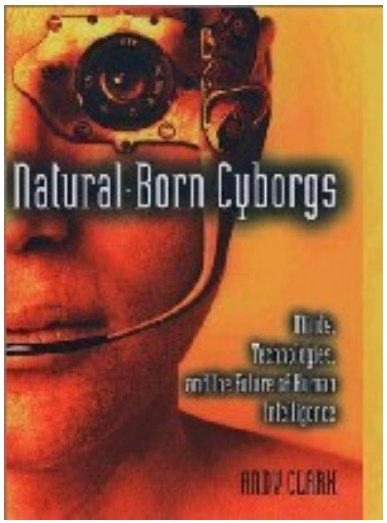

- Andy Clark, Oxford University Press, 2003
- "The drive toward biotechnological merger is deep within us"

Natural-born cyborgs?

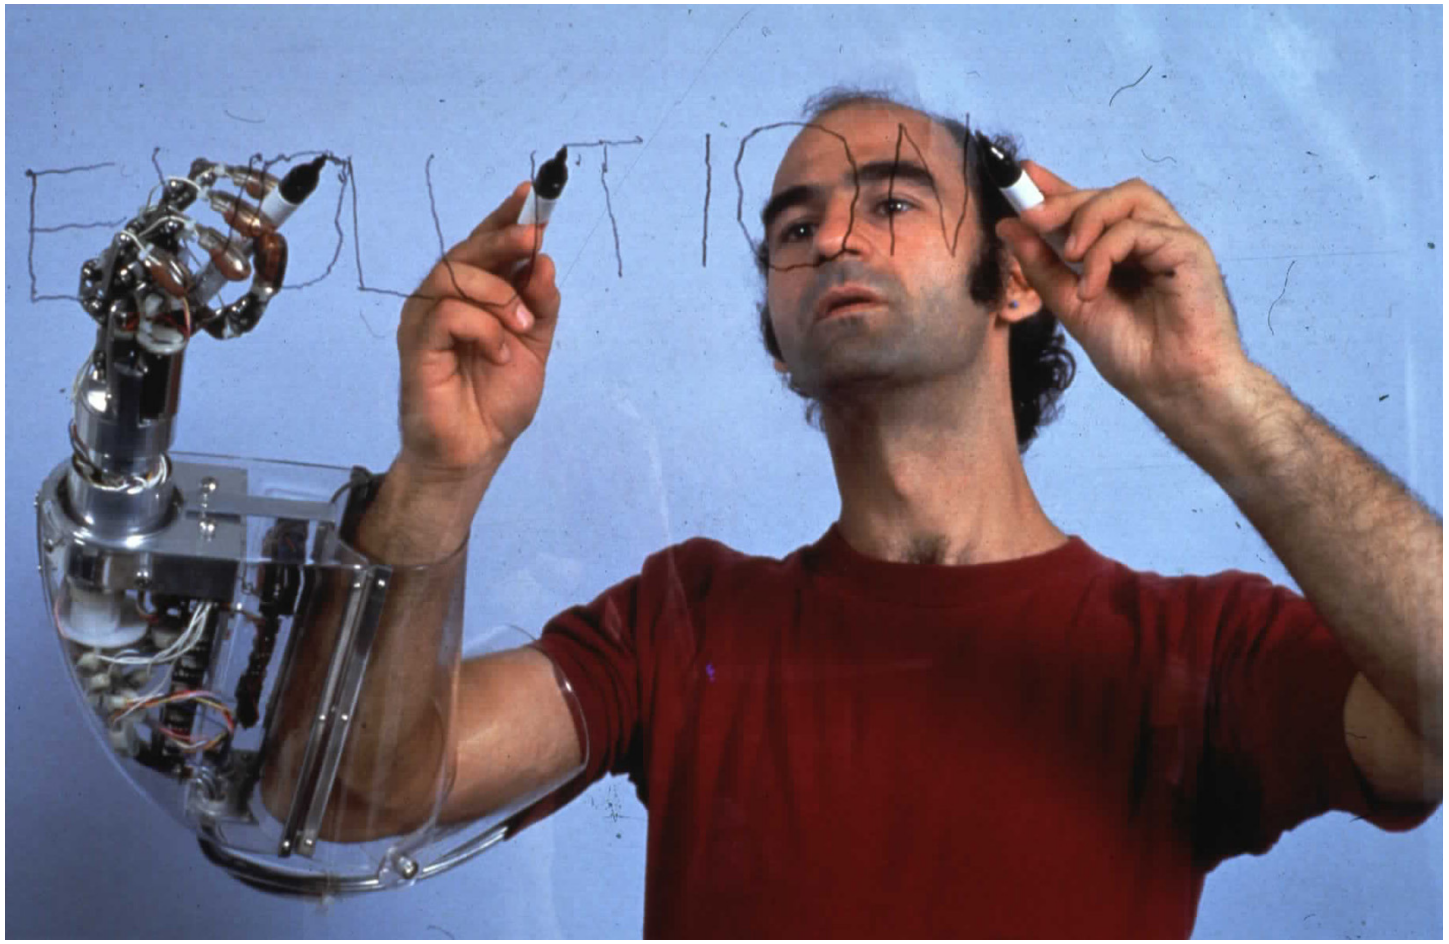

Mental prostheses

- While physical prostheses will play a role, mental prostheses will be much more widespread.
  - Ways to boost our knowledge/memory/processing capabilities
- No radical surgery necessary.
  - Though that is possible for those who want it)
- In fact we already rely on some simple mental prostheses

Mental Prostheses

- How do you answer the question:

  "Do you know what the time is?"

- How do you answer the question:

  "Do you know what *hymeneal* means?"

- How do you answer the question:

  "What is the best way to get to Bushwick?"

- Where next?

Next we have. . .


AT&T
9:42 AM
SMS
Text
9
Calendar
Photos
Camera
YouTube
Stocks
Maps
Weather
Clock
Calculator
Notes
Settings


htc
with Google
vodafone
10:10 AM
Dialer
Contacts
Browser
Vodafone
menu

Physical agents I

Physical agents II

The Two Key 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?

- These are the *micro* and *macro* perspectives.

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 give us a way to develop interaction-based systems.

Some Views of the Field II

- *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.

  Multi-agent approaches give us a way to simulate, and hence understand, different kinds of complex system.

Some Views of the Field III

- *Agents are the achievable bit of the AI project*:

  The aim of Artificial Intelligence as a field is to produce general human-level intelligence. This requires a very high level of performance in lots of areas:

  – Vision
  – Natural language understanding/generation
  – Reasoning

  Building an agent that can perform well on a narrowly defined task in a specific environment is much much easier (though not easy). Systems like DS1 show that this is possible.

Summary

- This has been a brief introduction to “An Introduction to Multiagent Systems”
- We have argued that MAS are:
  - a natural development of computer science;
  - a natural means to handle ever more distributed systems; and
  - not science fiction :-)
- We also made a first definition of “agent” and “multiagent system”.
- We will delve more into these things in the next lecture.