Artifacts for Agents: Multi-disciplinary Foundation

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna a Cesena

Academic Year 2013/2014

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology
   - Lessons Learned: From (Cognitive) Anthropology & Ethology to MAS

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Multi-, Inter- Trans-Disciplinary Research

From multi- to trans-disciplinary research

- Multi-disciplinary research means that multiple areas are involved in the same research activity—results are drawn from and concern different fields
- Inter-disciplinary research means that models, methods and techniques are brought from one area to a different one—results mainly concern the latter area
- Trans-disciplinary research means that models, methods and techniques are first brought from one area to a new one; then, once are suitably extended and generalised, results are brought back to the original area

Why X-Disciplinary Research? I

Convergence of Scientific Research

- Complexity of systems (observed, modelled, constructed) is characterising more or less all of the human knowledge
- The same patterns in observable phenomena, system structure & behaviour, scientific models, methods, and techniques, occur repeatedly in many heterogeneous research fields

Why X-Disciplinary Research? II

Convergence towards MAS

- Complexity of computational systems today matches complexity of biological, social, economical, organisational, . . . , systems
- Results from other areas dealing with complex systems may be useful / important / essential for computational systems & MAS in particular
- Results from computational systems & MAS are already changing the way in which scientific activity is conducted in every other areas

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Questions to be Answered I

We already learned something...

- ... about the reasons behind the agent abstraction,
- as well as some of its features

However, before a complete and precise definition could be given, some issues have to be clarified

Questions to be Answered II

We have to understand...

- ...if agents are the next thing after objects, what happens to objects, then? What about the paradigm shift?
- ...as object-oriented systems are made of interacting objects, are multiagent systems made of interacting agents—only?
- ...if societies and environment are essential to agent-oriented systems, how should they be handled in MAS modelling and engineering?
- ...if agents have to act, which are the objects of their acting?

Finally, we would like to taste...

- ...the flavour of X-disciplinary research

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Activity Theory (AT) I

Origins of (Cultural-Historical) Activity Theory

- Born in the context of Soviet Psychology
- Rooted in the dialectic materialism by Marx & Engels
- Mostly by the work by Lev Vygotsky (1926-62) [Vygotskiĭ, 1978]
- Broadly speaking, AT is a very general framework for conceptualising human activities – how people learn, how society evolves – based on the concept of human *activity* as the fundamental unit of analysis

Activity Theory (AT) II

Activity Theory nowadays

- Re-discovered and widely applied in Computer Science and related fields in the last years [Nardi, 1996]
- Mostly in fields like Computer Supported Cooperative Work (CSCW) and Human Computer Interaction (HCI)
- Brought to the MAS field by both Italian and Spanish groups—e.g. [Ricci et al., 2003]

Human Activity in AT

Main Focus of AT

- AT focuses on *human activities*
- within a *social / organisational context*
- as separated by their respective (physical and ideal) *objects*

Collaborative activities in AT

- Cooperation is understood as a *collaborative* activity
- A collaborative activity has *one objective*
- A collaborative activity is *distributed* onto *several actors*, who participate to the activity
- Explicit norms and rules regulate the relationships among individual participants' work

Mediated Interaction in AT

Every Human Activity...

- ...is found to be *mediated*...
- ...by mediating *artifacts*...
- ...of heterogeneous nature, either physical or psychological
  - operating procedures, heuristics, scripts, languages, ...

Artifacts in AT

- Artifacts are the tools that mediate actions and social interactions
  - artifacts mediate between individual participants and their *environment*
  - artifacts embody the portion of the environment that can be designed and controlled to support participants' activities
- As an observable part of the environment, artifacts can be monitored along with the development of the activities
  - to evaluate overall system performance and
  - to keep track of system history

Role of Artifacts in AT

- Artifacts can be either *physical* or *cognitive*—or, they may have a twofold nature
  - example of physical artifacts are shelves, doors, phones, whiteboards, ...
  - example of cognitive artifacts are operating procedures, heuristics, scripts, languages, ...
  - examples of artifacts with a twofold nature (physical / cognitive) are operating manuals, computers, ...
- Artifacts are both a *means* but also a *product* of social activity, so they embody a set of social practise
  - their design and structure reflect a history of particular use in some given social / organisational context

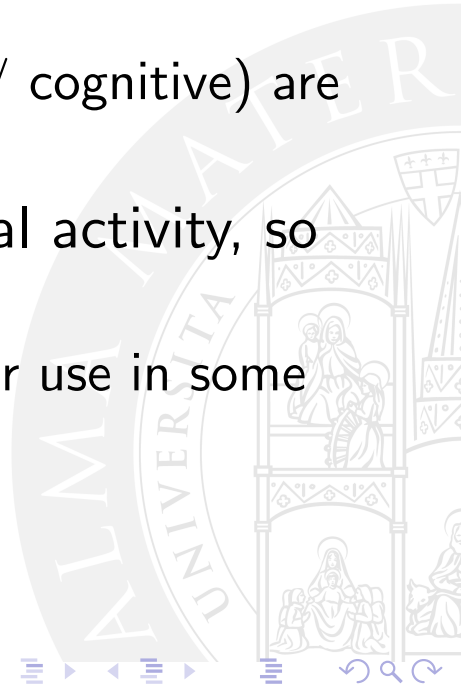

Artifacts as Enablers and Constrainers of Activities

- As mediating tools, artifacts have both an *enabling* and a *constraining* function

enablers artifacts expand out possibilities to manipulate and transform different objects

constrainers the object is perceived and manipulated through the artifact not 'as such' but within the limitations set by the artifact itself

- A simple example: a driving wheel

enabler enables me to change direction while driving a car

constrainer allows me only one way to change direction while driving a car

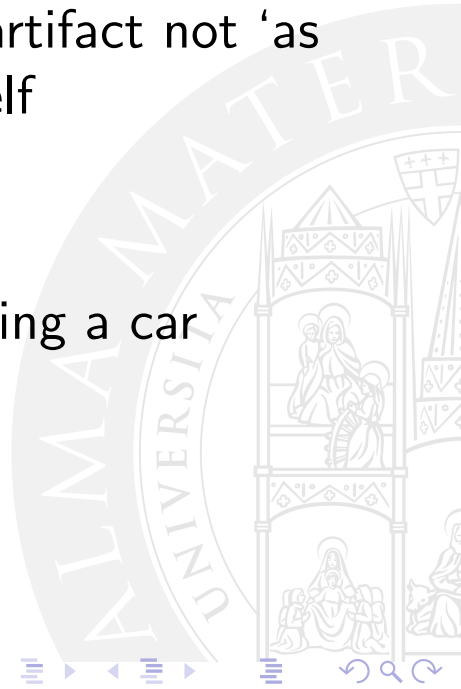

Layers for Collaboration Activities in AT I

- AT identifies a three-layered structure for social (collaborative) activities [Bardram, 1998, Engeström et al., 1997]
- The three layers are labelled as
  - *co-ordinated*
  - *co-operative*
  - *co-constructive*

Layers for Collaboration Activities in AT II

AT Layers: The Picture

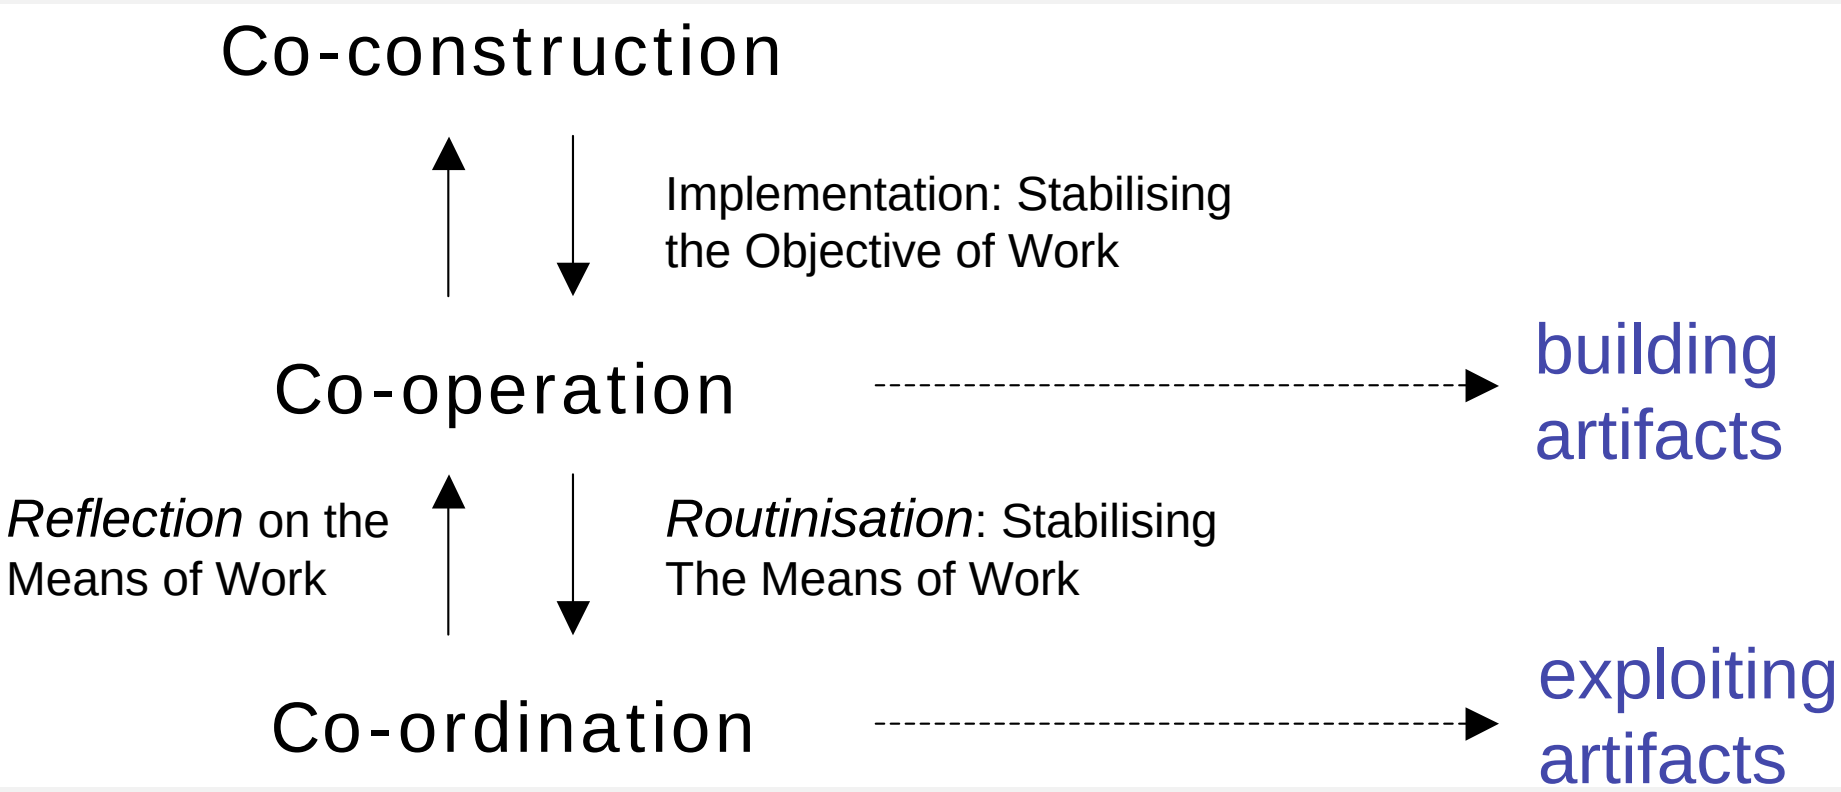

Co-ordination in AT

- The *co-ordinated* aspect of work captures the *normal* and routine *flow of interaction*
- Participants follow their *scripted roles*, each focusing on the successful performance of their actions, implicitly or explicitly assigned to them
- Participants share and act upon a common object, but their individual actions are only externally related to each other
- *Scripts* coordinating participants' actions are not questioned or discussed, neither known and understood in all their complexity
- Participants act as "wheels in the organisational machinery" [Kuutti, 1991], and co-ordination ensures that an activity is working in harmony with surrounding activities

Co-operation in AT

- The *co-operative* aspect of work concerns the mode of interactions in which actors focus on a common object, thus *share* the *objective* of the activity
- Here, actors do not have actions or roles explicitly assigned to them
- With regard to the common object, each actor has to balance his/her own actions with other agent actions, possibly influencing them to achieve the common task
- At the co-operation level
  - the object of the activity is stable and agreed upon
  - the means for realising the activity is not yet defined
- The means for realising a collaborative activity—the artifacts—are then the *object* of the co-operative activity, and its results as well

Co-construction in AT

- The *co-constructive* aspect of work concerns interactions in which actors focus on re-conceptualising their own organisation and interaction in relation to their shared objects
- Neither the object of work, nor the scripts are stable, and must be collectively constructed, i.e., *co*-constructed

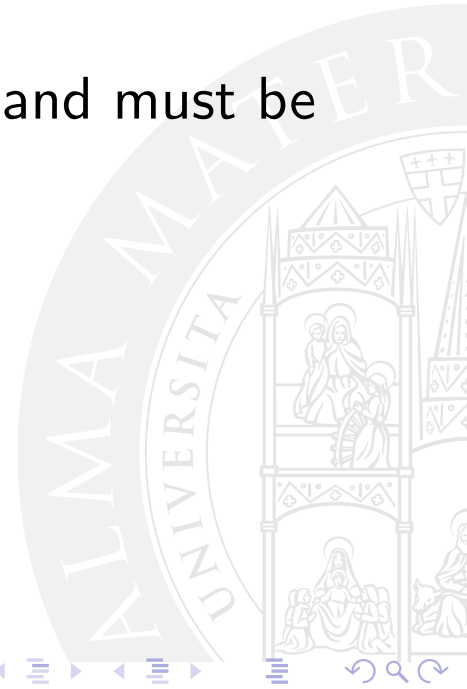

AT Layers: Summing Up

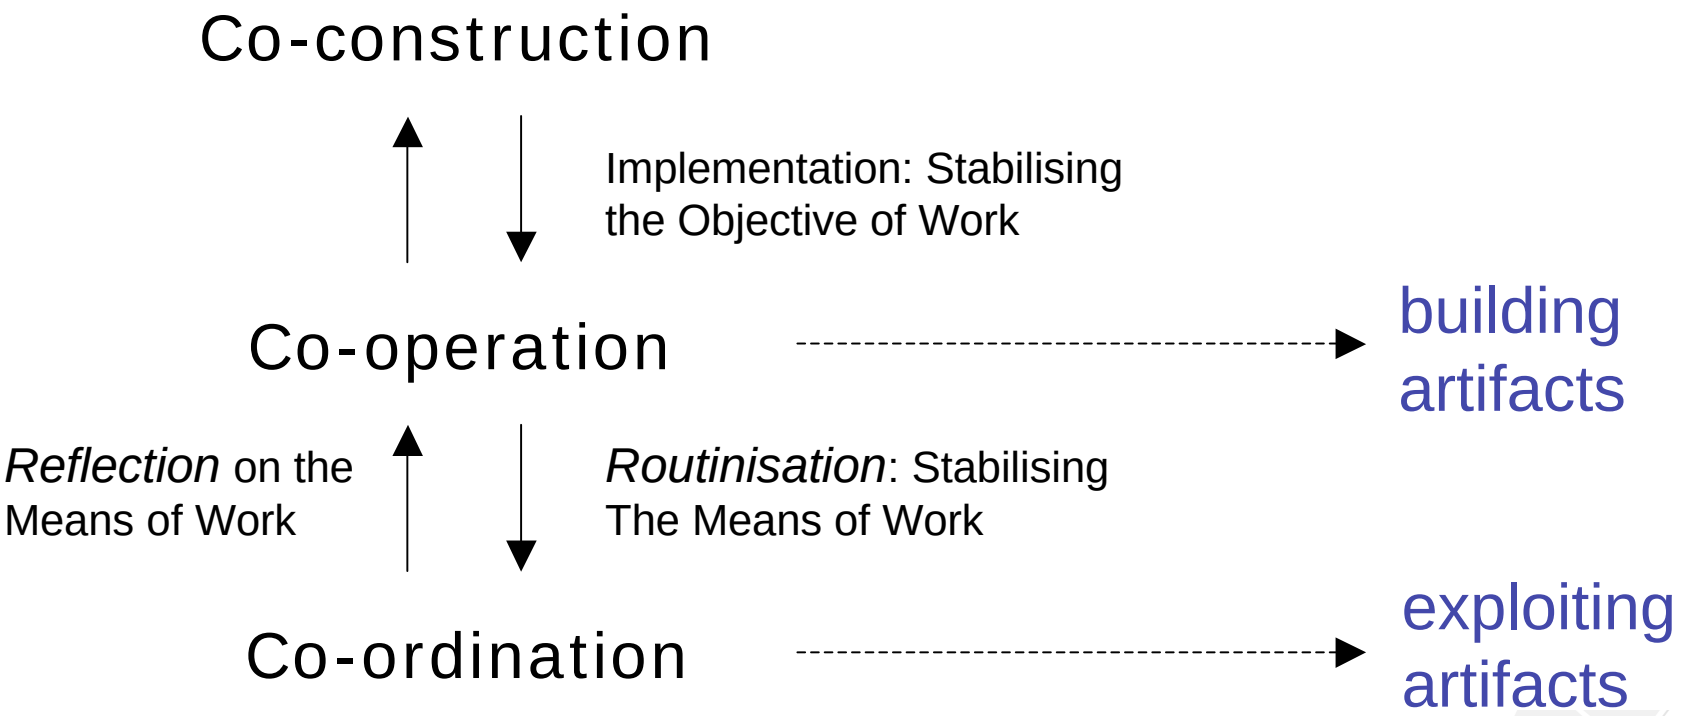

Collaborative activities in AT

- A collaborative activity is not to be seen in general at one single level
- Co-ordination, co-operation, and co-construction are instead to be interpreted as *analytical* distinctions of the same collaborative activity, concurring in different times and modes to its development

Outline

1 Premises
- Trans-disciplinary Research
- Dangling Issues

2 Activity Theory
- Background from Activity Theory
- Lessons Learned: From AT to MAS

3 Distributed Cognition
- Background from Distributed Cognition
- Lessons Learned: From Distributed Cognition to MAS

4 Sociology
- Background from Sociology
- Lessons Learned: From Sociology to MAS

5 Computer Supported Cooperative Work
- Background from CSCW
- Lessons Learned: From CSCW to MAS

6 (Cognitive) Anthropology & Ethology
- Background from (Cognitive) Anthropology & Ethology

Agents are not the Only Abstractions Needed

Basic Abstractions: Agents plus Artifacts

- Adopting AT as a conceptual framework for MAS social activities has led to recognise that *agents are not the only basic abstractions* to model and build MAS [Ricci et al., 2003]
- Artifacts, too, are necessary [Ricci et al., 2006]
  - to enable and constrain agent actions
  - to mediate agent interactions with other agents and with the environment
  - to model and shape MAS *environment*
  - in general, to improve agent ability to achieve their individual and social goals

Relevance of AT Research in MAS

Artifacts are essential—in MAS, too

- AT investigation is relevant in MAS since it points out that artifacts are essential to enable and govern agent actions and interactions within a MAS
  - by enhancing agent capabilities to act
  - by constraining both individual and social activities in a MAS

Role of environment

- AT emphasises the fundamental role of the *environment* in the development of complex systems
- Also, AT suggests that artifacts are the essential tools [Weyns et al., 2007, Viroli et al., 2005]
  - to model MAS environment
  - to shape it so as to make it favourable to the development of collaborative activities

Coordination Artifacts

Artifacts for collaboration and coordination

- *Coordination artifacts* are artifacts used in the context of collaborative activities, mediating the interaction among actors involved in the same social context [Ricci et al., 2003]
- Coordination artifacts can be either *embodied* or *disembodied*, referring to respectively physically or cognitive/psychological artifacts
- Coordination artifacts are social artifacts shared by agents in a MAS, which are meant to enable and govern the interaction among agents, and between agents and their environment

Coordination artifacts & media

- Coordination artifacts represent a straightforward generalisation of the notion of coordination medium, as coming from fields like coordination models and languages and distributed AI
- Examples include abstractions like tuple spaces, channels, blackboards, but also pheromone infrastructures, e-institutions, ...

AT Layers for MAS Collaboration

Layers for MAS collaboration & coordination artifacts

- The three levels identified by AT for social activities can be re-interpreted in the MAS context in terms of the relationship between agents and artifacts—in particular, coordination artifacts
- The three layers are labelled as
  - *co-ordination*
  - *co-operation*
  - *co-construction*

AT Layers for MAS in Detail

co-construction — agents understand and reason about the (social) objectives (goals) of the MAS, and build up a model of the social tasks required to achieve them—this also involves identifying interdependencies and interactions to be faced and managed

co-operation — agents design and build the coordination artifacts—either embodied (coordination media) or disembodied (plans, interaction protocols, etc.)—which are useful to carry on the social tasks and to manage the interdependencies and interactions devised out at the previous (co-construction) stage

co-ordination — agents use the coordination artifacts: then, the activities meant at managing interdependencies and interactions—either designed a-priori or planned at the co-operation stage—are enforced/automated

Levels of Use of Artifacts

Co-ordination: both intelligent and non-intelligent agents could coordinate

Any agent (either intelligent or not) can simply exploit artifacts to achieve its own goals by simply taking artifacts as they are, and use them

Co-operation: intelligent agents could change artifacts to change MAS

Intelligent agents could possibly reason about the nature of the artifacts as well as on the level of achievement of their goals, and take the chance to change or adapt the artifacts, or even to create new ones whenever useful and possible as the result of either an individual or a social activity

Co-operation: MAS engineers could embody social intelligence in artifacts

In the same way, MAS engineers can use artifacts to embody the "social intelligence" that actually characterises the systemic/synergistic (as opposed to compositional) vision of MAS [Ciancarini et al., 2000], but also to observe, control, and possibly change MAS social behaviour

AT Layers for MAS Collaboration: The Picture

AT Layers for MAS: The Picture

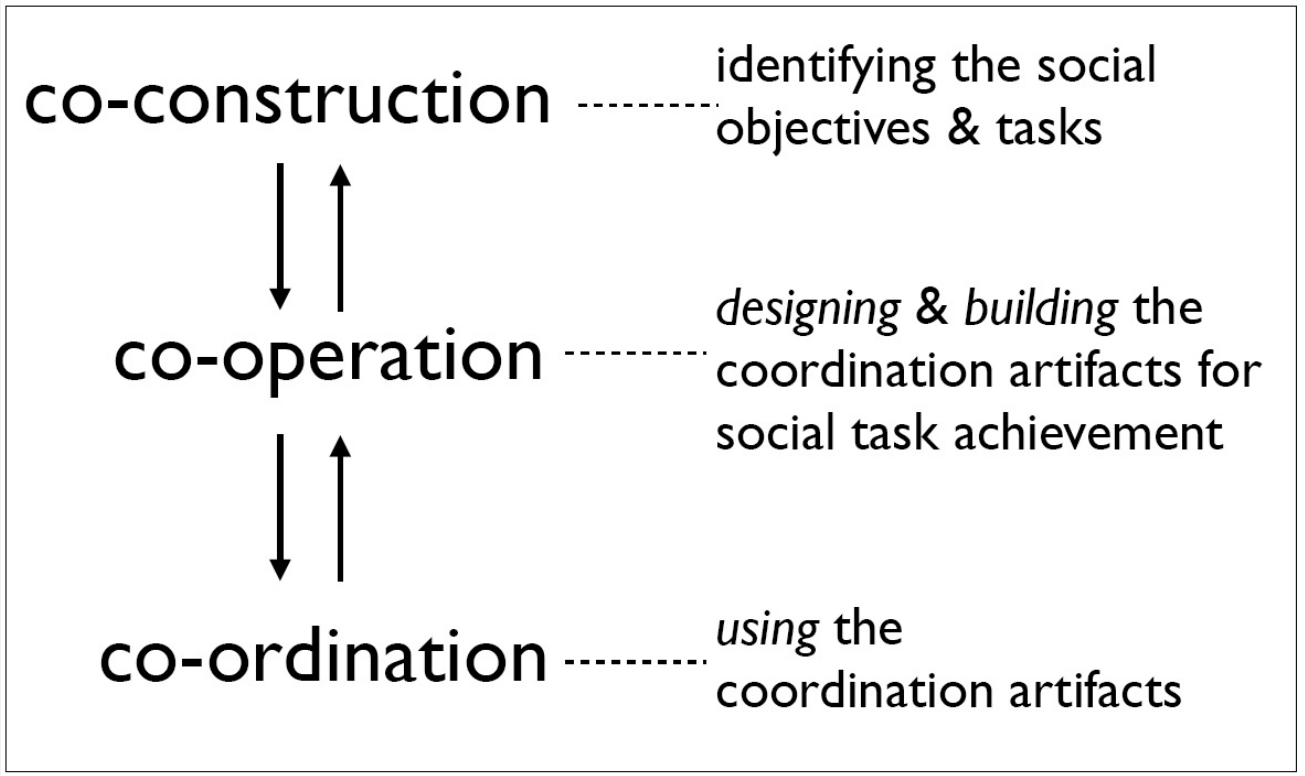

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Distributed Cognition in Short

Distributed Cognition. . . [Kirsh, 1999]

- . . . is a branch of cognitive sciences
- which proposes that *human cognition* and *knowledge representation*
  - rather than solely confined within individuals
  - is *distributed* across individuals, tools and *artifacts* in the *environment*

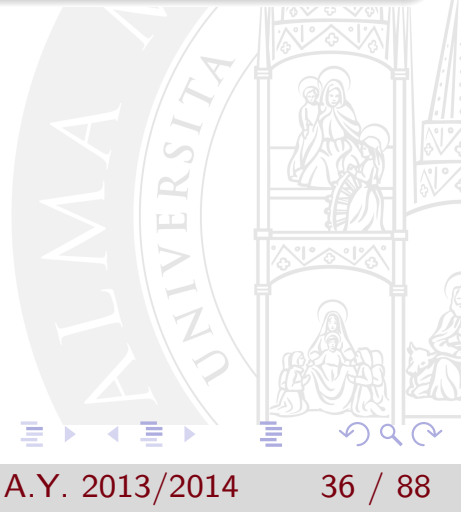

Cognition is Distributed

Cognition transcends individuals

- Intelligent processes in human activity go beyond the boundaries of individual actors
- Knowledge is not confined within human minds
- Cognition transcends individual cognition

Knowledge representation transcends individuals

- Knowledge representation does not pertain individual humans only
- Representation is distributed
  - partially in the mental spaces of humans
  - as external representations of memories, facts, and information of any sort distributed on the objects, tools and instruments that constitute the environment

Distributed Cognitive Systems

Analysis of Distributed Cognition focuses on *distributed cognitive systems*

- People interact with external *cognitive artifacts* containing knowledge represented in some form
- Human *intelligent behaviour* results from the distributed interactions with other humans and with cognitive artifacts
- In the overall, this defines and determines the context where human activities are *situated*
  - that is, the physical, cultural and social context that also guides, constrains and partially determines intelligent activities

Cognitive Artifacts

Cognitive artifacts: a definition [Norman, 1992]

those artificial devices that maintain, display, or operate upon information in order to serve a representational function and that affect human cognitive performance

Cognitive artifacts are...

- ...a product of human design and work
- ...aimed at aiding or enhancing our cognitive abilities
  - like post-its, calendars, agendas, computers, etc.
- ...not mere amplifiers of our cognitive abilities
  - cognitive artifacts also modify the nature of the tasks to be performed

Personal vs. System View

System view

Individuals plus artifacts altogether as a (functional) subsystems

- Understanding activities requires to consider (cognitive) actors and (cognitive) artifacts altogether
- Actions are sometimes mediated sometimes targeted to artifacts, and cannot be fully understood without them

Personal view

Individuals as subsystems affected by artifacts

- Practical reasoning is deeply affected by artifacts
- Individuals should change the way in which they represent actions, plan, deliberate and finally act

Environment in Distributed Cognitive Systems

Environment has a key role in distributed cognitive systems

In distributed cognitive systems, the nature of the environment

- on the one hand, depends on the artifacts and tools that shape it
- on the other hand, determines the efficiency and effectiveness of the work and activities of the actors that are immersed in it

Work environment

- How do we define a working environment for individuals and organisations?
  - it mostly depends on the tasks that have to be carried on inside
- Real work environments are a complex superposition of social, cultural, cognitive, and physical constraints
- How should the environment be understood as a complex analytical construct when the goal is environment design?

Coordination in Distributed Cognitive Systems

Observing real world activities

- An effective environment for a successful activity is a shifting coalition of resources and constraints
  - some physical, some social, some cultural, some computational
  - involving both internal and external computational resources
- Activity is successful whenever such a coalition is suitably coordinated
  - lack of coordination means failure of activity
- Coordination is then essential, and concerns activities, resources and constraints
  - at both the individual and the social level

The Function of Action in Distributed Cognitive Systems

What is the purpose of an activity?

- A dominant assumption is that the point of activity is to change the environment in a way that (presumably) leads to goal satisfaction
- Many action however do not make sense under this assumption
  - most communication actions, but not only them
  - for instance, people undertake actions to save attention, memory and computation; people recruit external elements to reduce their own cognitive effort by distributing computational load
  - this make sense if people is situated
- As a result, environment design should not merely be aimed at helping people to achieve their goals
  - it should also be designed to make other actions easy
  - such as epistemic, complementary, coordinative actions

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Cognition Outside Agents

Cognition & knowledge representation do not belong to agents only

- Objects & tools in the environment may participate to the cognitive processes
- Structure of MAS environment may explicitly represent knowledge

Cognition & knowledge representation are distributed in the environment

- Artifacts are essential parts of the MAS cognitive processes
- Cognitive artifacts encapsulate knowledge as explicitly represented

Agent View vs. MAS View

Personal / agent view

- Once artifacts are exploited, they change the way in which agents act and reason about action

System / MAS view

- In order to understand and possibly evaluate agent (social) action within a MAS, one should consider agent(s)+artifact(s) altogether

MAS Environment is Structured

(Cognitive) artifacts shape MAS environment

- Artifacts determine the structure of MAS environment
- Knowledge is distributed in the environment, and encapsulated within cognitive artifacts
- Structure of the environment, and knowledge it contains, affect the activities of agents within MAS

MAS Action & Coordination

MAS coordination depends on environment structure

- Environment structure changes the nature of agent action
- Environment structure affects agent mutual interaction
- Environment structure modifies the way agents coordinate in a MAS
- Environment structure should be designed to
  - help agent actions to achieve their goals
  - help epistemic, complementary, coordinative agent actions easier / effective

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Agents & Goals I

Cognitive interpretation of (social) action [Conte and Castelfranchi, 1995]

- Agents in a society can be generally conceived as either *goal-governed* or *goal-oriented* entities
  - goal-governed entities refer to the strong notion of agency, i.e. agents with some forms of cognitive capabilities, which make it possible to explicitly represent their goals, driving the selection of agent actions
  - goal-oriented entities refer to the weak notion of agency, i.e. agents whose behaviour is directly designed and programmed to achieve some goal, which is not explicitly represented
- In both cases, agent goals are *internal*

Agents & Goals II

External goals

- *External goals* refer to goals that typically belong of the social context or *environment* where the agents are situated
- External goals are sorts of regulatory states which condition agent behaviour
  - a goal-governed system follows external goals by adjusting internal ones

Entities without Goals

Not every entity involved in (social) actions has a goal

- Within a society, there are entities that are explicitly designed to provide a *function*
- *Artifacts* are such objects
  - they have a function associated
- Artifacts have no goals to achieve
  - they may have a *destination* associated
    - a *destination* is external goal attached to an artifact by an agent, in the act of using it
  - destination is then associated to the *use* of an artifact
    - destination is related but not identical to function: an artifact can be used according to a destination that differs from its designed function

On the Relation Between Agents & Artifacts [Conte and Castelfranchi, 1995]

Use & use value

- When facing an artifact, an agent may adopt different perspectives
- Evaluating an artifact for use, to select it among many others, and then to use it, to achieve agent's own goals, are two different matters
- Different sorts of external goals are associated by an agent to an artifact

use value the *use-value* goal, according to which the artifact should allow user agents to achieve their objective—this drives the agent *selection* of the artifact

use the *use* goal, which directly corresponds to the agent internal goal—this guides the actual *usage* of the artifact

Articulating the Agents & Artifacts Relationship

How could an agent deal with an artifact?

- There are at least three different ways an agent can exploit an artifact

  use by merely using it, according to its function, and associating it to a destination

  selection by selecting it for future use, according to its function, its possible future destinations, and the agent's goals and plans

  construction & manipulation by adapting & changing an existing artifact, or by creating a new one for future use, thus designing its function, according to its possible future destinations, and the agent's goals and plans

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Goals in MAS

Agents have goals

strong agency Agents have explicitly-represented goals

weak agency Agents have implicitly-represented / encoded goals

Artifacts have functions

- Artifacts have no internal goals
- Artifacts have a pre-designed function
- An artifact is associated with an external goal (its destination) by agents in the act of using it

Agents & Artifacts Interacting

Aspects of agent-artifact relationship

use An agent can use an artifact, according to its use goal, associating it with a destination

- **aware use** because the agent is aware of the artifact's function
- **unaware use** because the artifact's use is encoded in the agent by the programmer / designer

selection An agent can select an artifact for future use, according to its use-value goal, reasoning about its possible future destinations and use goals

construction / manipulation An agent can modify an artifact to adapt its function to some required use-value goals and to its possible future destinations

- or, an agent can create *ex-novo* a new artifact with an agent-designed function according to some required use-value goals and to its possible future destinations

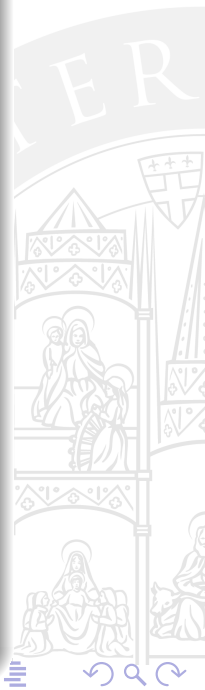

MAS Engineers Designing Agents

Basic choices to make in agent design

- Should an agent be aware of artifact's behaviour and structure, and of how to use them?
  - should an agent be able to reason and deliberate about artifact use?
- Should an agent be aware of artifact's function and possible uses?
- Should an agent be able to act over artifacts to modify them and adapt their function?
  - should an agent be able to create *ad hoc* artifacts *ex novo*?
- Should a MAS engineer be able to act over artifacts to modify them and adapt their function, or, to create new artifacts, *at run-time*?

MAS Engineers Designing Artifacts

Basic issues in artifact design

- How should an artifact be made in order to be ready for agent's use?
  - either aware, or unaware
  - possibly, within an open system
- How should an artifact be made in order to be ready for agent's evaluation and selection?
- How should an artifact be made in order to be ready for agent's modification and adaptation?
- How should MAS environment be structured in order to allow artifact run-time creation and modification?
  - by agents and MAS engineers?

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Computer Supported Cooperative Work (CSCW)

Basic issues in CSCW

- CSCW aims at automating human cooperative work through computational procedures
- However, two diverging strategies are currently emerging [Schmidt and Simone, 2000]

  automation stressing computational procedures to automate coordination of activities

  flexibility stressing the flexibility of computational procedures with respect to intelligent coordination by collaborating actors

- The former approach emphasises coordination by the computational entities ruling collaboration, the latter coordination by intelligent collaboration entities
- Main problem: the two strategies *diverge*, they should instead *converge*

Automation vs. Flexibility: Key Issues in CSCW I

Mutual awareness for flexibility

- *Mutual awareness* means that the actors of a collaboration activity affect and mutually perceive the other actor's activities through the shared workspace
- The so-called *common field of work* can reveal / conceal portions of the collaboration activities to the participants
- Mutual awareness is then the basis for opportunistic, *ad hoc* alignment and improvisation, which ensure *flexibility* to collaborative activities

Automation vs. Flexibility: Key Issues in CSCW II

Coordinative artifacts for automation

- *Coordinative artifacts* are the rulers of collaboration
- They work more as *constrainers* rather than as *commanders*
- By giving structure to the common field of work, coordinative artifacts encapsulate those coordination responsibilities that are better to be *automatised* in order to achieve efficiency in cooperation
- In all, coordinative artifacts

  work as constrainers they define and govern the space of the admissible articulation of activities

  work *not* as commanders they do not impose a pre-defined course of actions that could cause unnecessary rigidity and reduce the required flexibility

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Automation of Collaboration Activities in MAS

Coordinative artifacts for automation of MAS collaboration

- *Coordinative artifacts* rule MAS collaboration, working more as constrainers rather than as commanders
- Coordinative artifacts structure MAS common field of work, as specialised abstractions automatising and making collaboration efficient
- As constrainers, coordinative artifacts define and govern the space of the admissible articulation of MAS collaboration activities
- On the other hand, they do not impose a pre-defined course of actions, promoting flexibility of intelligent agent coordination, and respecting agent autonomy

Flexibility of Collaboration Activities in MAS

Mutual awareness for flexibility of MAS collaboration

- Shared MAS environment should be structured as the MAS *common field of work* to allow agents to mutually perceive each other's activities (*mutual awareness*)
- MAS common field of work can reveal / conceal portions of MAS collaboration activities to the agents
- Mutual awareness promotes opportunistic alignment and improvisation of agent activities, and ensure *flexibility* to MAS collaboration

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

Anthropology: The Logocentric Philosophical Bias I

Human capacity of language as the main sign of intelligence?

- Western anthropology has long dwelt on such a point
- The relation between language, use of tools, and evolution of intelligence has long been neglected [Hewes, 1993]

Anthropology: The Logocentric Philosophical Bias II

Human capacity of developing and using tools as a fundamental sign of intelligence

- Humans forge & use tools
- The first characterisation of *Homo Abilis* is its ability to forge tools
  - tool using vs. tool making distinction
- This is a clear sign of intelligence
- Evidence of co-evolution of language and tools use along with human intelligence is overwhelming in modern anthropological studies [Gibson and Ingold, 1993]

Ethology: Tools, Animals, and the Tooling Test

Use of tools is not an exclusive feature of humans

- Beavers build dams, bees construct perfect hexagonal cells, many birds live in self-made nests, ...
- Ethologists commonly measure intelligence of animals by making them face problems that require the use of tools to be solved (e.g. [Povinelli, 2000])
- A sort of tool-equivalent of the Turing test has been proposed by philosopher Ronald Endicott, which was aimed at evaluating intelligence in terms of the ability to exploit tools
  - the so-called "Tooling Test for Intelligence" [Wood et al., 2005]

Anthropology from Theologists: Tools & Self-Awareness

Use of tools reveals awareness of self [Martelet, 1998]

- When using a tool, a creature shows it is able to distinguish and identify itself from the world around
- The use of a tool *reveals awareness* of self, and of the environment as well
  - whenever a tool is built with a goal, it is stored for further / repeated use, it is used for building new tools, etc.
- Tools are at the same time the first and the most distinctive *expression of human intelligence*, along with language
- They are the most powerful amplifiers of the (both individual and social) *human ability to affect the environment*—to survive environment change, first, and to change the environment for the human purposes, then

Outline

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology

The Logocentric Philosophical Bias in MAS I

Agent capacity of language as the main sign of agent intelligence?

- Research on MAS still dwells on the *logocentric bias*
- Intelligent *use of tools by agents* is typically neglected
  - as a stunning example, FIPA (Foundation for Intelligent Physical Agents) just ignore pragmatic / physical agent actions, and only focuses on agent communication actions

The Logocentric Philosophical Bias in MAS II

Agent ability of developing and using tools as a sign of agent intelligence

- A notion of *tool* for agents is needed
- Agents should be *able to use tools*
- Intelligent agents should be *able to forge & adapt tools*
- A theory of physical / pragmatical action should be developed for agents, as refined as the one for communication actions
- Such a theory should focus on tool use / creation by agents
  - The notion of *Agens Faber* goes along this very direction [Omicini et al., 2006]

Tools, Agents, and the Tooling Test

Use of tools should be a feature for agents in a MAS

- MAS researchers should be able to measure intelligence of agents by making them face problems that require the use of tools to be solved
- A sort of tool-equivalent of the Turing test for agents using tools should be defined, aimed at evaluating agent intelligence in terms of the ability to exploit tools
  - a sort of "Tooling Test for Agent Intelligence"
- Agent intelligence should then be measured by both the agent ability to communicate and by agent ability to use tools
  - the two abilities should be somehow strictly related, and "co-evolve" in some sense—a common theory of agent action could be of use here

1. Premises
   - Trans-disciplinary Research
   - Dangling Issues
2. Activity Theory
   - Background from Activity Theory
   - Lessons Learned: From AT to MAS
3. Distributed Cognition
   - Background from Distributed Cognition
   - Lessons Learned: From Distributed Cognition to MAS
4. Sociology
   - Background from Sociology
   - Lessons Learned: From Sociology to MAS
5. Computer Supported Cooperative Work
   - Background from CSCW
   - Lessons Learned: From CSCW to MAS
6. (Cognitive) Anthropology & Ethology
   - Background from (Cognitive) Anthropology & Ethology
   - Lessons Learned: From (Cognitive) Anthropology & Ethology to MAS

References I

Bardram, J. (1998).
Designing for the dynamics of cooperative work activities.
In *1998 ACM Conference on Computer Supported Cooperative Work (CSCW'98)*, pages 89–98. ACM Press.

Ciancarini, P., Omicini, A., and Zambonelli, F. (2000).
Multiagent system engineering: The coordination viewpoint.
In Jennings, N. R. and Lespérance, Y., editors, *Intelligent Agents VI. Agent Theories, Architectures, and Languages*, volume 1757 of *LNAI*, pages 250–259. Springer-Verlag.
6th International Workshop (ATAL'99), Orlando, FL, USA, 15–17 July 1999. Proceedings.

References II

Conte, R. and Castelfranchi, C., editors (1995).
Cognitive and Social Action.
Routledge.

Engeström, Y., Brown, K., Christopher, L. C., and Gregory, J. (1997).
Coordination, cooperation, and communication in the Courts: Expansive transitions in legal work.
In Cole, M., Engeström, Y., and Vasquez, O. A., editors, *Mind, Culture, and Activity. Seminal Papers from the Laboratory of Comparative Human Cognition*, chapter 28, pages 369–388.
Cambridge University Press.

Gibson, K. R. and Ingold, T., editors (1993).
Tools, Language and Cognition in Human Evolution.
Cambridge University Press.

References III

Hewes, G. W. (1993).
A history of speculation on the relation between tools and languages.
In [Gibson and Ingold, 1993], pages 20–31.

Kirsh, D. (1999).
Distributed cognition, coordination and environment design.
In Bagnara, S., editor, *3rd European Conference on Cognitive Science (ECCS'99)*, pages 1–11, Certosa di Pontignano, Siena, Italy. Istituto di Psicologia, Consiglio Nazionale delle Ricerche.

Kuutti, K. (1991).
The concept of activity as a basic unit of analysis for CSCW research.
In Bannon, L. J., Robinson, M., and Schmidt, K., editors, *2nd European Conference on CSCW (ECSCW'91)*, pages 249–264. Kluwer Academic Publisher.

References IV

Martelet, G. (1998).
Évolution et création, tome 1 – Sens ou non-sens de l'homme dans la nature?
Editions du Cerf, Paris.

Nardi, B. A., editor (1996).
Context and Consciousness: Activity Theory and Human-Computer Interaction.
MIT Press.

Norman, D. A. (1992).
Design principles for cognitive artifacts.
Research in Engineering Design, 4(1):43–50.

References V

Omicini, A., Ricci, A., and Viroli, M. (2006).
Agens Faber: Toward a theory of artefacts for MAS.
Electronic Notes in Theoretical Computer Sciences, 150(3):21–36.
1st International Workshop "Coordination and Organization" (CoOrg 2005), COORDINATION 2005, Namur, Belgium, 22 April 2005. Proceedings.

Povinelli, D. J. (2000).
Folk Physics for Apes: The Chimpanzee's Theory of How the World Works.
Oxford University Press.

References VI

Ricci, A., Omicini, A., and Denti, E. (2003).
Activity Theory as a framework for MAS coordination.
In Petta, P., Tolksdorf, R., and Zambonelli, F., editors, *Engineering Societies in the Agents World III*, volume 2577 of *LNCS*, pages 96–110. Springer-Verlag.
3rd International Workshop (ESAW 2002), Madrid, Spain, 16–17 September 2002. Revised Papers.

Ricci, A., Viroli, M., and Omicini, A. (2006).
Programming MAS with artifacts.
In Bordini, R. P., Dastani, M., Dix, J., and El Fallah Seghrouchni, A., editors, *Programming Multi-Agent Systems*, volume 3862 of *LNAI*, pages 206–221. Springer.
3rd International Workshop (PROMAS 2005), AAMAS 2005, Utrecht, The Netherlands, 26 July 2005. Revised and Invited Papers.

References VII

Schmidt, K. and Simone, C. (2000).
Mind the gap! Towards a unified view of CSCW.
In Dieng, R., Giboin, A., Karsenty, L., and De Michelis, G., editors, *Designing Cooperative Systems: The Use of Theories and Models*, volume 58 of *Frontiers in Artificial Intelligence and Applications*, Sophia Antipolis, France. IOS Press.
4th International Conference on the Design of Cooperative Systems (COOP 2000), Proceedings.

Viroli, M., Omicini, A., and Ricci, A. (2005).
Engineering MAS environment with artifacts.
In Weyns, D., Parunak, H. V. D., and Michel, F., editors, *2nd International Workshop "Environments for Multi-Agent Systems" (E4MAS 2005)*, pages 62–77, AAMAS 2005, Utrecht, The Netherlands.

References VIII

Vygotskiĭ, L. S. (1978).
Mind in Society: Development of Higher Psychological Processes.
Harvard University Press.

Weyns, D., Omicini, A., and Odell, J. (2007).
Environment as a first-class abstraction in multi-agent systems.
Autonomous Agents and Multi-Agent Systems, 14(1):5–30.
Special Issue on Environments for Multi-agent Systems.

Wood, A. B., Horton, T. E., and Amant, R. S. (2005).
Effective tool use in a habile agent.
In Bass, E. J., editor, *2005 IEEE Systems and Information Engineering Design Symposium (SEADS 2005)*, pages 75–81, Charlottesville, VA, USA. IEEE.

Artifacts for Agents: Multi-disciplinary Foundation

Distributed Systems
Sistemi Distribuiti

Andrea Omicini
andrea.omicini@unibo.it

Dipartimento di Informatica – Scienza e Ingegneria (DISI)
Alma Mater Studiorum – Università di Bologna a Cesena

Academic Year 2013/2014