

10th European Agent Systems Summer School – EASSS'08 –

Mehdi Dastani João Leite (eds.)

Faculdade de Ciências e Tecnologia
Universidade Nova de Lisboa
2008

Preface

The European Agent Systems Summer School (EASSS) aims to offer a valuable forum for knowledge exchange between various research groups in its field for the benefit of students and researchers at both beginner and advanced level. Attracting about 100 participants each year, EASSS provides a wide range of state-of-the-art courses which are broad enough to provide a general introduction to the chosen topic, whilst also covering the most important contributions in depth. The success of the summer school series has been essential not only for the mutual benefit of research groups, but also in forming future generations of researchers to maintain the excellence of Agents research and development in Europe.

The 10th edition of EASSS will take place at the Science and Technology Faculty of the New University of Lisbon, located on the south bank of the Tagus River, near the beautiful beaches of Caparica, between the 5th and the 9th of May, 2008. Previous editions of EASSS took place in Durham, UK (2007), Annecy, France (2006), Utrecht, Netherlands (2005), Liverpool, UK (2004), Bologna, Italy (2003), Barcelona, Spain (2002), Prague, Czech Republic (2001), Saarbrücken, Germany (2000) and Utrecht, Netherlands (1999).

The 10th Edition of EASSS features 16 courses, invited or selected by the EASSS Advisory Board from a pool of submitted proposals, covering a very wide range of topics. Follows their abstracts:

Introduction to Multiagent Systems by M. Luck

This course will provide an introduction to the field of agents and multi-agent systems, spanning both the micro and macro level, including: basic concepts; agent architecture; interaction and multi-agent systems; applications; and future challenges.

Logics for Multiagent Systems by J.-J. Ch. Meyer

In this tutorial I will treat the basics of modal logic, and how they are applied to describe agent attitudes and multi-agent systems. Topics that will be addressed are basic modal logic, epistemic logic, temporal logic, dynamic logic, deontic logic as well as BDI logic, Common Knowledge and Joint Intentions, logic as in applied in normative systems.

Agent Swarms Generating Short-term Forecasts and Increasing Situational Awareness by P. Valckenaers

In application areas such as logistics, traffic and manufacturing, agent technology is employed to perform coordination. This requires a look-ahead capability from the agents. In a distributed setting, this look-ahead capability has to account for interactions amongst the participants. Mainstream negotiation-based agent technologies perform poorly in this matter. Two designs based on swarms of light-weight agents offer a solution. The course compares two approaches for propagating the intentions of situated agents through a virtual world, mirroring some part of the real world. This propagation is executed by swarms of lightweight agents. The propagation of agent intentions causes short-term forecasts to emerge. In a first approach, the world-of-interest is mirrored in a "book" in which each page is a picture of the situation. Each page corresponds to an instance in time, either the present, recent past or imminent future. The state and the interactions on page n determine the prediction of the system state on page $n+1$. In the second approach, the world-of-interest is mirrored in a graph-shaped information structure. Each node representing a resource (e.g. a road segment, truck, machine, etc.) has an agenda. In this agenda, users make reservations that reflect their intentions. The intelligent resources answer queries about their future services whilst accounting for this agenda. This second approach follows a travel arrangement pattern. Users explore the graph-shaped world model through virtual execution of a possible journey. Having explored sufficiently, the users select a best-performing journey to become their intention, and repeat the virtual execution of this journey whilst entering reservations at the resources involved.

Trust and Reputation in Multiagent Systems by G. Muller and L. Vercouter

The course is intended for people that want to get a wide perspective on computational trust and reputations models. The tutorial will be mainly composed of three parts. The first part presents the general motivations and problems tackled by trust and reputation models. It distinguishes these problems from the complementary ones tackled by security techniques and poses the foundations of trust models, based on sociological works about trust and reputation. The second part consists in a state of the art of the existing works in the field of computational trust and reputation in multiagent systems. It addresses prominent models proposed in this domain: it starts from simple models like Marsh's trust model or eBay and OnSale reputation model, then present more decentralized models like Schillo and Funk's model and ends with very rich and social models like ReGret and LIAR. The last part focuses on the ART (Agent Reputation and Trust) testbed. It is an open platform for experimentation and competition for trust and reputation models. Its aims at providing means to compare such models. The end of this last part of the tutorial is dedicated to practice on the ART platform.

Service Oriented Agents by B. Hirsch

One of the current hypes in the IT industry is the application of service oriented architectures (SOA) to a wide range of software systems that are distributed and often cross organisational boundaries. It also appears that most of the promises of SOA are

also made by agent technology - be it interoperability or agility - yet it is webservices and SOAP instead of agents and speechacts that is en vogue today. It is however difficult to discern between marketing hype and actual contributions to the application of distributed systems. The aim of this tutorial is to present the field of services and service oriented architectures (SOA), and to show in how far agents and SOA and related technologies can benefit from each other. We believe that this topic is very relevant to aspiring agent researchers as service orientation is of great importance in the industry, and as the aims of agent research and SOA are similar. We therefore want to create awareness of the implications of SOA and the potential benefits the two areas can have from each other. The tutorial will consist of the following three parts: Introduction to SOA: Students will be given a short introduction to service orientation and SOA in general. The pros and cons of SOA will be presented to, as well as practiced by the students using handouts and small presentations. How agents can contribute to Services and vice versa: In the second part, we will focus on the common ground of agents and services. We will point out different shortcomings of the current SOA approach as well as possible solutions that agents provide. Possible topics here are (formal) semantics in general, ontologies, coordination and autonomous planning, goal oriented behaviour and matchmaking. We also present possible lessons that agent research can learn from SOA and associated technologies. Current Approaches: The final part of the tutorial deals with work at the interface between agents and services. In particular, we plan to present our own agent framework JIAC, as well as work to incorporate (web-)services, BPEL, and agents.

Introduction to Game Theory and Mechanism Design by P. Harrenstein and M. de Weerd

In multi-agent systems multiple agents with individual preferences interact in a common environment. For the design and proper formal understanding of the interaction in multi-agent systems appropriate mathematical concepts are required. Game theory provides and analyzes such concepts as equilibrium outcomes and strategies. Roughly speaking, mechanism design is concerned with the development of multi-agent systems with desirable game-theoretic properties even if the preferences of the agents are unknown to the designer or to the other agents. It is also concerned with what is formally possible and impossible in this respect. Put this way, the relevance of game theory and mechanism design for multi-agent research might seem obvious. Nevertheless, game theory nor mechanism design are commonly part of a standard curriculum in computer science. This course aims to fill this gap. This introductory course is aimed at post-graduate students in computer science and multi-agent systems who have no previous acquaintance with game theory or mechanism design. As such only some elementary mathematical training is assumed. The main goal of the course is to make the students familiar with the most fundamental concepts and results of game theory and mechanism design and to make clear how game theory and mechanism design can be applied to their own research.

Wireless Sensor Networks and Multiagent Systems by E. Platon and D. Weyns

Wireless sensor networks (WSN) are networks of small computers endowed with sensors and wireless connectivity. The comparison of WSN with multi-agent systems (MAS) is natural, as the two approaches consist of interactive entities situated in an environment they can sense and act upon locally. Agents provide engineers with a higher level of abstraction, so that WSN become a useful application domain for MAS. MAS researchers have actually developed adequate models and approaches to address some current research issues in the area of WSN (e.g. self-adaptation and collective behaviors). The goals of this tutorial are to (i) introduce WSN technologies, (ii) show why MAS are important for the development of WSN, (iii) explain the main research challenges in WSN, and (iv) show how MAS can address relevant issues. This tutorial is a half-day session tailored for researchers and practitioners interested in the application of MAS to WSN. It assumes background knowledge in MAS and basics of networks, although it will summarize important information when necessary. Related events: Workshop on Agent Technologies for Sensor Networks (AAMAS); International Conference on Networked Sensing Systems; Conference on Embedded Networked Sensor Systems; European Conference on Wireless Sensor Networks, RoboCup.

Foundations of Institutions by M. De Vos, M. Esteva and J. Padget

Institutions are multi-agent systems in which agents are governed by norms. Agents can break the norms provided that they are willing to accept the ensuing sanction. With agents not constantly abiding by the rules, it becomes important that institutions can monitor and penalise any violations that occur. In this six hour tutorial we will:

- Discuss the basic principles of institutions and their components;
- Present a general framework that allows a user to specify, design, reasoning about and verify institutional models;
- Describe and demonstrate graphical notations and tools for the creation and animation of institutions.

The primary delivery method will be by lecture alleviated by class discussion and exercises.

Agents and Arguments by S. Modgil and L. Amgoud

An argument is a set of premises offered in support of a conclusion, and argumentation is the process whereby arguments are constructed, exchanged and evaluated in light of their interactions with other arguments. Logical models of argumentation have recently emerged as a promising paradigm for modeling agent reasoning and communication. The paradigm's promise resides in its modular nature that closely mirrors the way humans reason. It thus provides a general framework for inference and decision making over the full gamut of mental attitudes (beliefs, goals, desires, actions e.t.c) in the presence of inconsistent, uncertain and incomplete information. The generality of the argumentation paradigm is substantiated by the fact that existing logical approaches to non-monotonic and commonsense reasoning can be formalised in terms of argumentation. Furthermore, argumentation has recently been used to enrich models of communication. It provides a principled way in which to structure rational dialogue and

enable exchange of, and reasoning about, justifications/arguments for proposals and or statements between human and or automated agents. This course aims to introduce students and researchers to the fundamentals of logic based models of argumentation, and to present recent research work on development of these models for application in agent reasoning and communication. We review: 1) Abstract argumentation frameworks and extensions to these frameworks to facilitate flexible and adaptive agent reasoning. 2) Instantiations of these abstract frameworks for formalising agent inference and decision making over mental attitudes. 3) Frameworks for argumentation based dialogue, whereby agents argue in order to persuade or negotiate with other agents.

Computational Complexity in Multiagent Systems by E. Elkind and E. Markakis

In many multi-agent settings, agents have limited computational resources, and cannot function in environments where each decision requires large amounts of computation. Therefore, computational complexity is an important parameter in designing and evaluating multiagent systems. It provides a useful tool to the agent designer, enabling him to choose whether to focus on finding an optimal strategy, or use approximate or heuristic approaches. In this tutorial, we cover foundations on computational complexity and its applications in the context of multiagent systems. We first introduce some of the most important complexity classes, such as P, NP, PSPACE and #P, as well as the basic notions of computational complexity, such as reductions, and explain what it means for a problem to be hard or complete for a given class. We then provide examples of multiagent decision problems that are hard for the above mentioned classes, and discuss the implications of these results from a practical perspective. Usually, a hardness result motivates research in heuristic and approximation algorithms for the problem, and we provide several examples where such approaches succeed. On the other hand, we also demonstrate that computational hardness can be used as a barrier to undesirable behavior by strategic agents, in settings where a truthful agent's decision problem is easy, while finding a way to manipulate the system to an agent's benefit is computationally hard. Our examples will be drawn from various areas of research in multiagent systems, such as multiagent planning, argumentation, auctions, coalition formation and voting, illustrating the importance of computational complexity concepts in these domains. The course is intended for 1st-year PhD and MSc students with little or no background in computational complexity.

Planning in Multiagent Systems by M. de Weerd and C. Witteveen

By definition, agents are autonomous entities that are able to act. Hence, the process of determining which actions to execute and the order in which these actions have to be executed is considered as an essential property of agents. Within the AI-community such processes have been studied in the context of planning problems. Traditionally, the planning community concentrated on planning as a single-agent phenomenon, focusing on aspects of representation, efficiency and robustness of a plan executed by a single agent. In a multi-agent context, however, other aspects of the planning process

become more important. Nowadays a major issue in multi-agent planning is the coordination of cooperative as well as self-interested agents. The focus on coordination in planning also brings in other major topics in multi-agent systems like auctions, coalition formation, privacy, scalability, etc. In this tutorial we start with an overview of multi-agent planning problems and discuss the properties of these problems that have the most influence on the applicability of certain existing approaches. We will give an overview of these approaches by distinguishing their role in the planning and coordination phases of a multi-agent system. We then discuss how single agent AI planning techniques lead to different multi-agent planning approaches. Finally, we discuss some particular approaches to coordination of single-agent planning systems in more detail.

What Coalitions Can Achieve by J. Dix and W. Jamroga

In this course, we present some formal approaches to modeling and reasoning about strategic abilities of groups of agents. The course consists of two parts. The first part addresses prescriptive concepts (mostly from collaborative game theory) that specify how much a collaboration between agents is worth and what it means for a team to be "good". We also discuss various search algorithms to find such good coalitions (or, ones within a certain bound from the optimal coalition). The focus of the second part is descriptive. We introduce a modal logic of strategic ability, ATL, and show how abilities of coalitions can be specified in that language. Finally, we discuss the algorithmic side of checking such specifications. The course requires some elementary knowledge of logic and game theory. Familiarity with basic modal logic will be an advantage. Preliminary outline of the course: 1. How to Form a Coalition. - Basic models & concepts (brief overview): strategic games, extensive games, from extensive to strategic form, games with imperfect information; Nash equilibrium, Pareto optimality. - Coalition formation: definition of the problem, core, CS-search, payoff division (Shapley value). - Contract Nets. 2. Reasoning about What Teams Can Achieve. - Modal logic (brief intro): models, operators, epistemic logic. - Reasoning about abilities of agents and teams: alternating-time temporal logic ATL for perfect information games, problems with imperfect information, ATLir. Model checking abilities of coalitions. - Reasoning about rational teams.

Agent Oriented Software Engineering by O. Shehory

Agent Oriented Software Engineering (AOSE) is a key factor for introducing agent-based systems to the industry as an engineering approach. At present, the majority of existing agent applications are developed in an ad hoc fashion: little or no rigorous design methodology, limited specification of the requirements, ad hoc design of agents and of multi-agent systems as a whole, and little attention to non-functional requirements such as mobility, scalability, performance issues, and standards. By adopting AOSE principles, one gains the advantages of an organized development process such as reusability, testing, and maintenance. One of the basic principles of AOSE is using a methodology for developing agent applications. Hence, this course will concentrate on methodologies,

their applicability, and their use. In particular, the goals of this course are the following: -Introduce basic concepts of software engineering in the context of agent-based systems; -Introduce the motivation for using agent-oriented software engineering; -Introduce the field of agent-oriented methodologies; -Present several agent-oriented methodologies; -Compare existing agent-oriented methodologies; -Discuss implementation issues of agents and MAS and their relationship to agent-oriented methodologies.

Automated Negotiations in Electronic Markets by N. Gatti

Agents are well-suited for dynamic, constrained, and real-time environments such as electronic marketplaces. In such environments agents representing their users negotiate for goods and services following negotiation protocols. Bargaining and auctions are the principal negotiation protocols for buying and selling goods based upon competition among the interested parties. This tutorial will introduce participants to agent-based negotiations. The tutorial will start by introducing agent-based negotiation and negotiation protocols in general. Bargaining and auctions will then be described in detail. Essential concepts that are required for following the tutorial will be introduced along the way. The bargaining problem will be introduced and the principal non-cooperative bargaining protocol (Rubinstein's alternating-offers) will be discussed in detail. Variations and applications Rubinstein's protocol in computer science will be presented. The four single side auction protocols (English, Dutch, FPSB, Vickrey) will be discussed in depth along with their relative advantages and disadvantages. Double auctions and the M-th and (M+1)-st clearing rules will also be covered. A brief exposition into more advanced auction formats such as multi-attribute and combinatorial auctions will follow.

Agent Based Simulation for Social Studies by L. Antunes and F. Cecconi

Agent-based social simulation is a recent multi-disciplinary effort that has increasingly established new challenges for the agents community, by bringing the agent technology to face complex phenomena such as the ones found in social sciences. At the same time, social scientists have been discovering how the computer and especially the advances in the agent field can provide a new and exciting tool to tackle the problems of their field, providing a paradigm shift in social sciences. The exchange between researchers in both areas has proven mutually fruitful, as much inspiration in MAS has come from social sciences. The course begins with an introduction to multi agent modelling: the main issues are the micro-macro links and the concept of rationally-heterogeneous agents. Then, the tutorial will deal with the main concepts of game theory and dynamical systems, following this outline: an overview of game theory and equilibrium (Nash, repeated games, Bayesian games); learning and evolutionary games (ESS); dynamical systems; Monte-Carlo and numerical simulations. Finally, agent based methodology will be discussed: (a) the problem of the level of description, (b) the relationships with mathematical modelling, (c) the tools and the problem of scalability. Using NetLogo, students will carry out some experiments: (a) cooperation in minority

games; (b) reputation's dynamic; (c) social capital in a social networks; (d) imitation vs. evolution in an artificial trade market.

Normative Multiagent Systems by G. Boella, D. Grossi and L. van der Torre

The Agentlink Roadmap, published in 2005, considers norms as key for the development of MAS. Applications of norms range from agent organizations and electronic institutions (e-commerce, and e-government), to open agent societies, agent communication, trust and reputation systems, and MAS programming. Norms are among the social notions which are obtaining most attention in the MAS research area (see, for instance, the COIN and NorMAS workshops, the last editions of the DEON conference), bringing together MAS researchers, legal and social scientists, logicians. The tutorial provides a clear understanding about when norms become important in designing and developing MAS, about what kind of norms are of importance for MAS and how they should be used. A student session is planned on the application of norms to the students' research projects. The tutorial is given by leading scientists in the field of Normative Multi-Agent Systems, organizers of the NorMAS'05-'07-'08 workshops and of the DEON'08 conference. The tutorial is structured in a session of three lectures and a student session: 1. When are norms to be used?: The First part of the tutorial provides an introduction to Normative Multi-Agent Systems stressing what are the MAS-related topics in which norms play a key role; 2. What kind of norms are to be used?: The second part of the tutorial focuses on the different types of norms (and their interaction) which are of particular relevance for the design of Normative Multi-Agent Systems; 3. How are norms to be used?: The third part of the tutorial shows how norms can be used to support the design of several aspects of Normative Multi-Agent Systems and how norms relate to other social notions such as roles, organizations and institutions; 4. Student session: Students will be actively involved by discussing relations of their research questions to the topics dealt with in the tutorial.

The 10th Edition of EASSS is sponsored by the International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), Fundação para a Ciência e Tecnologia (FCT), Center for Artificial Intelligence of UNL (CENTRIA) and Associação Portuguesa para a Inteligência Artificial (APPIA).

This book contains the course notes of the 10th European Agent Systems Summer School.

April 2008

Mehdi Dastani
João Leite

EASSS'08 Organisation

EASSS'08 Committee

João Leite (Local Chair), New University of Lisbon, Portugal

Mehdi Dastani (Chair), Utrecht University, The Netherlands

Rafael H. Bordini, University of Durham, UK

Catholijn Jonker, Delft University, The Netherlands

Barbara Keplicz, Warsaw University, Poland

Michal Pechoucek, CTU Prague, Czech Republic

Wiebe van der Hoek, University of Liverpool, UK

Gerhard Weiss, Software Competence Center Hagenberg, Austria

Local Organisation

João Leite (Local Chair), New University of Lisbon, Portugal

Martin Slota, New University of Lisbon, Portugal

Sponsoring Institutions



FCT Fundação para a Ciência e a Tecnologia

MINISTÉRIO DA CIÊNCIA, TECNOLOGIA E ENSINO SUPERIOR

