



The Seventh International Conference on Autonomous Agents and Multiagent Systems

Estoril, Portugal

May 12-16, 2008

Abstracts

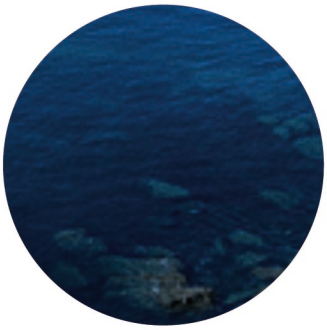





AAMAS 2008 - Abstracts

This document contains all abstracts. It is organized as follows: first all abstracts of full papers organized by sessions and tracks, then all abstracts of short papers organized by topic. All papers are (additionally in case of full papers) presented in form of a poster. Therefore every abstract is marked with a number plus a letter. This letter helps you to find the relevant poster session:

- posters A is 4.30-5.20,
- posters B is 5.20-6.10 and
- posters C is 6.10-7pm

All poster sessions will take place on Wednesday May 14. For the rooms of each poster session, please consult the program booklet.

... means that the paper has been nominated for a best paper award (AAMAS 2008 Best Paper Award, Best Student Paper Award or Best Industry and Application Track Award).

Wednesday, May 14th, 8:40-9:40, Invited Talk

Autonomous Virtual Humans and Lower Animals: From Biomechanics to Intelligence

Demetri Terzouloulos

The confluence of virtual reality and artificial life, an emerging discipline that spans the computational and biological sciences, has yielded synthetic worlds inhabited by realistic artificial flora and fauna. Artificial animals are complex synthetic organisms that have functional, biomechanical bodies, perceptual sensors, and brains with locomotion, perception, behavior, learning, and cognition centers. Virtual humans and lower animals are of interest in computer graphics because they are self-animating graphical characters poised to dramatically advance the interactive game and motion picture industries even more so than have physics-based simulation technologies. More broadly, these biomimetic autonomous agents in their realistic virtual worlds also foster deeper computationally-oriented insights into natural living systems. Furthermore, they engender interesting applications in computer vision, sensor networks, archaeology, and other domains.

Wednesday, May 14th, 10:10-12:10 – Session 1

Virtual Agents I

Chair: Jean-Claude Martin

1A

An Empathic Virtual Dialog Agent to Improve Human-Machine Interaction

Magalie Ochs, Catherine Pelachaud, David Sadek

Recent research has shown that virtual agents expressing empathic emotions toward users have the potentiality to enhance human-machine interaction. To identify under which circumstances a virtual agent should express empathic emotions, we have analyzed real human-machine dialog situations that have led users to express emotion. The results of this empirical study have been combined with theoretical descriptions of emotions to construct a model of empathic emotions. Based on this model, a module of emotions has been implemented as a plug-in for JSA agents. It determines the empathic emotions (their types and intensity) of such agents in real time. It has been used to develop a demonstrator where users can interact with an empathic dialog agent to obtain information on their emails. An evaluation of this agent has enabled us to both validate the proposed model of empathic emotions and highlight the positive user's perception of the virtual agent.

4A

ERIC: A Generic Rule-based Framework for an Affective Embodied Commentary Agent

Martin Strauss, Michael Kipp

We present ERIC, an affective embodied agent for realtime commentary in many domains. The underlying architecture is rule-based, generic, and lightweight -- based on Java/Jess modules. Apart from reasoning about dynamically changing events, the system can produce coherent natural language and nonverbal behaviour, based on a layered model of affect (personality, mood, emotion). We show how reasoning, template-based natural language generation and affective appraisal can be implemented within the same rule-based paradigm. To make the system domain-independent we worked on two different domains, a virtual horse race and a multiplayer tank battle game. We empirically evaluated the genericness of the system by measuring the effort it takes to change the domain, and discuss the results.

7A

The Identification of Users by Relational Agents

Daniel Schulman, Mayur Sharma, Timothy Bickmore

Virtual agents designed to establish relationships with more than one user must be able to identify and distinguish among those users with high reliability. We describe an approach for relational agents in public spaces to identify repeat users based on two strategies: a biometric identification system based on hand geometry, and an identification dialogue that references previous conversations. The ability to re-identify visitors enables the use of persistent dialogue and relationship models, with which the agent can perform a range of behaviors to establish social bonds with users and enhance user engagement. The agent's dialogue encourages users towards repeat visits, and provides mechanisms of recovery from identification errors, as well as contextual information which may be used to improve the accuracy of the biometric identification. We have implemented and evaluated this identification system in a virtual guide agent for a science museum that is designed to conduct repeated and continuing interactions with visitors. We also present the results of a preliminary evaluation of the system, including user opinions of this technology, and of the effect of identification, both successful and unsuccessful, on acceptance and engagement of the agent.

10A

Trackside DEIRA: A Dynamic Engaging Intelligent Reporter Agent

François Knoppel, Almer Tigelaar, Danny Oude Bos, Thijs Alofs, Zsofi Ruttkay

DEIRA is a virtual agent commenting on virtual horse races in real time. DEIRA analyses the state of the race, acts emotionally and comments about the situation in a believable and engaging way, using synthesized speech and facial expressions. In this paper we discuss the challenges, explain the computational models for the cognitive, emotional and communicative behavior, and account on implementation and feedback from users.

13A

Does the Contingency of Agents' Nonverbal Feedback Affect Users' Social Anxiety?

Sin-Hwa Kang, Jonathan Gratch, Ning Wang, James Watt

We explored the association between users' social anxiety and the interactional fidelity of a virtual human, specifically addressing whether the contingency of agents' nonverbal feedback affects the relationship between users' social anxiety and their feelings of rapport, performance, or judgment on interaction partners. This subject was examined across four experimental conditions where participants interacted with three different types of agents and a real human. The three types of agents included the Non-Contingent Agent, the Responsive Agent (opposite to the Non-Contingent Agent), and the Mediated agent (controlled by a real human). The results indicated that people having greater anxiety would feel less rapport and show worse performance while feeling more embarrassment if they experience the untimely feedback of the Non-Contingent Agent. The results also showed people having more anxiety would trust real humans less as their interaction partners. We discuss the implication of this relationship between social anxiety in a human subject and the interactional fidelity of a virtual human on the design of virtual characters for social skills training and therapy.

16A

The Design of a Generic Framework for Integrating ECA Components

Hung-Hsuan Huang, Aleksandra Cerekovic, Igor Pandzic, Yukiko Nakano, Toyoaki Nishida

Embodied Conversational Agents (ECAs) are life-like computer generated characters that interact with human users in face-to-face multi-modal conversations. ECA systems are generally complex and difficult for individual research groups to develop. If there was a common framework for connecting ECA functioning blocks seamlessly to an integrated system, many redundant efforts can be saved. This paper discusses the issues emerged in developing such framework and introduces the design and preliminary results of our ongoing project, Generic ECA Framework that is composed with a blackboard based communication platform, an XML based high-level protocol and a set of API libraries.

Agent-based System Development I

Chair: Carles Sierra

57C

Cost-Based BDI Plan Selection for Change Propagation

Khanh Hoa Dam, Michael Winikoff

Software maintenance is responsible for as much as two thirds of the cost of any software, and is consequently an important research area. In this paper we focus on the change propagation problem: given a primary change that is made in order to meet a new or changed requirement, what additional, secondary, changes are needed? We build on previous work that has proposed to use a BDI (belief-desire-intention) agent framework to propagate changes by fixing violations of consistency constraints. One question that needs to be answered as part of this framework is how to select between different applicable (repair) plan instances to fix a given constraint violation? We address this issue by defining a suitable notion of repair plan cost that incorporates both conflict between plans, and synergies between plans. We then develop an algorithm, based on the notion of cost that finds cheapest options and proposes them to the user.

60C

Integrating Artifact-Based Environments with Heterogeneous Agent Programming Platforms

Alessandro Ricci, Michele Piunti, Lemi Daghan Acay, Rafael Bordini, Jomi Hübner, Mehdi Dastani

The "Agents and Artifacts" (A&A) and CARTAGO infrastructure are becoming increasingly popular as respectively a conceptual & programming model and related infrastructure for developing artifact-based environments, as are platforms based on agent-oriented programming languages. However, so far there has been no work on developing multi-agent systems where agents implemented and deployed in different agent-programming platforms can interact as part of the same multi-agent system, with a shared environment. Due to the generality of CARTAGO environments and its Java-based implementation, we have successfully implemented a multi-agent system where Jason, 2APL, (as BDI based approaches) and simpA (which is activity-oriented rather than BDI) agents work together by sharing an artifact-based environment. This paper shows how this was achieved by allowing Jason, 2APL, and simpA agents to interact with CARTAGO artifacts, describing a general approach for developing such heterogeneous multi-agent systems.

63C

A Domain Specific Modeling Language for Multiagent Systems

Christian Hahn

Software systems are becoming more and more complex with a large number of interacting partners often distributed over a network. A common dilemma faced by software engineers in building complex systems is the lack of clear requirements and domain knowledge needed to come up with a detailed design of the system. Agent technologies are a suitable programming paradigm to cope with the complexity of modern software systems. However, existing agent-based methodologies and tools are developed for experienced programmers and are not suitable for non-agent experts. This paper discusses a domain specific modeling language for multiagent systems that (i) provides a clear syntax and semantics to define agent-based systems in a graphical visualized manner and (ii) can be used to automatically derive code from its design through model transformations.

Agent Communication

Chair: Shaheen Fatima

70A

Constitutive Interoperability

Amit Chopra, Munindar Singh

In recent work, commitments have emerged as a valuable high-level notion for supporting autonomy in open

multiagent systems---particularly when the principals are business entities. Traditionally interoperation is approached from the standpoint of data exchange or of messaging. We use commitments to characterize interoperability in high-level terms: at the level of the communications among agents. Specifically, two agents are interoperable if their commitments align. Drawing upon Kant's famous distinction, we distinguish between two kinds interoperability, constitutive and regulative. Constitutive interoperability takes into account solely the meaning of messages whereas regulative interoperability also takes into consideration message order, occurrence, and data flow. We present a language for specifying agents constitutively and a decision procedure for determining their interoperability.

71B

Annotation and Matching of First-Class Agent Interaction Protocols

Tim Miller, Peter McBurney

Many practitioners view agent interaction protocols as rigid specifications that are defined a priori, and hard-code their agents with a set of protocols known at design time – an unnecessary restriction for intelligent and adaptive agents. To achieve the full potential of multi-agent systems, we believe that it is important that multi-agent interaction protocols are treated as first-class computational entities in systems. That is, they exist at runtime in systems as entities that can be referenced, inspected, composed, invoked and shared, rather than as abstractions that emerge from the behaviour of the participants. Using first-class protocols, a goal-directed agent can assess a library of protocols at runtime to determine which protocols best achieve a particular goal. In this paper, we present three methods for annotating} protocols with their outcomes, and matching protocols using these annotations so that an agent can quickly and correctly find the protocols in its library that achieve a given goal. We discuss the advantages and disadvantages of each of these methods.

****72C****

Agent Communication in Ubiquitous Computing: the Ubismart Approach

Jurriaan van Diggelen, Robbert-Jan Beun, Rogier van Eijk, Peter J. Werkhoven

This paper studies the use of agent communication in ubiquitous computing. This application domain allows us to investigate the efficient handling of large quantities of information in agent-based systems. We will present an approach to dynamically set up a communication network between agents which aims to minimize the communication load. The approach is based on a formal ontological notion of informativeness, on quantitative measures such as information gain and on the proper use of interaction mechanisms such as Publish/Subscribe. We also present experimental results which have been obtained using our prototyping tool called Ubismart.

73A

Using Organization Knowledge to Improve Routing Performance in Wireless Multi-Agent Networks

Huzaifa Zafar, Victor Lesser, Dan Corkill, Deepak Ganesan

Multi-agent systems benefit greatly from an organization design that guides agents in determining when to communicate, how often, with whom, with what priority, and so on. However, this same organization knowledge is not utilized by general-purpose wireless network routing algorithms normally used to support agent communication. We show that incorporating organization knowledge (otherwise available only to the application layer) in the network-layer routing algorithm increases bandwidth available at the application layer by as much as 35 percent. This increased bandwidth is especially important in communication-intensive application settings, such as agent-based sensor networks, where node failures and link dynamics make providing sufficient inter-agent communication especially challenging.

74B

Conjunctive Queries for Ontology based Agent Communication in MAS

Cassia Trojahn, Paulo Quaresma, Renata Vieira

In order to obtain semantic interoperability in open Multi-Agent Systems, agents need to agree on different ontologies. In this paper we formally define the notion of a mapping between ontologies, where mappings are expressed as correspondences between queries over ontologies. First, individual mappings are computed by specialized agents using different mapping approaches (lexical, semantic and structural). Next, these agents use argumentation to exchange their local results, in order to agree on the obtained mappings. An Extended Value-based Argumentation Framework (E-VAF) is used to represent arguments with *strength*. The E-VAF allows to determine which arguments are acceptable, with respect to the different *audiences* represented by different agents. To each argument is associated a *strength*, representing how confident an agent is in the similarity of two ontology terms. Based on their preferences and confidence of the arguments, the agents compute their preferred mapping sets. The arguments in such preferred sets are viewed as the set of globally acceptable arguments. These arguments are then represented as conjunctive queries in OWL-DL extended with DL-safe rules, a restriction imposed to attain decidability in such query answering system.

75C

A Cooperation-Based Approach For Evolution Of Service Ontologies

Murat Sensoy, Pinar Yolum

Communication among agents requires a common vocabulary to facilitate successful information exchange. One way to achieve this is to assume the existence of a common ontology among communicating agents. However, this is a strong assumption, because agents may have experiences that evolve their ontologies independently. When this is the case, agents need to form common grounds to enable communication. Accordingly, this paper proposes an approach in which agents can add new service concepts into their service ontologies and teach others services from their ontologies by exchanging service descriptions. This leads to a society of agents with different but overlapping ontologies where mutually accepted services emerge based on agents' exchange of service descriptions. Our simulations of societies show that allowing cooperative evolution of local service ontologies facilitates better representation for agents' needs. Further, through cooperation, not only more useful services emerge over time, but also ontologies of agents having similar service needs become aligned gradually.

Agent Societies and Societal Issues

Chair: Michael Luck

119B
Robust Normative Systems
Thomas Agotnes, Wiebe van der Hoek, Michael Wooldridge
Although normative systems, or social laws, have proved to be a highly influential approach to coordination in multi-agent systems, the issue of *compliance* to such normative systems remains problematic. In all real systems, it is possible that some members of an agent population will not comply with the rules of a normative system, even if it is in their interests to do so. It is therefore important to consider the extent to which a normative system is *robust*, i.e., the extent to which it remains effective even if some agents do not comply with it. We formalise and investigate three different notions of robustness and related decision problems. We begin by considering sets of agents whose compliance is necessary and/or sufficient to guarantee the effectiveness of a normative system; we then consider quantitative approaches to robustness, where we try to identify the proportion of an agent population that must comply in order to ensure success, and finally, we consider a more general approach, where we characterise the compliance conditions required for success as a logical formula.

122B
Sequential Decision Making with Untrustworthy Service Providers
Luke Teacy, Georgios Chalkiadakis, Alex Rogers, Nick Jennings
In this paper, we deal with the sequential decision making problem of agents operating in computational economies, where there is uncertainty regarding the trustworthiness of service providers populating the environment. Specifically, we propose a generic Bayesian trust model, and formulate the optimal Bayesian solution to the exploration-exploitation problem facing the agents when repeatedly interacting with others in such environments. We then present a computationally tractable Bayesian reinforcement learning algorithm to approximate that solution by taking into account the expected value of perfect information of an agent's actions. Our algorithm is shown to dramatically outperform all previous finalists of the international Agent Reputation and Trust (ART) competition, including the winner from both years the competition has been run.

125B
A Statistical Relational Model for Trust Learning
Achim Rettinger, Matthias Nickles, Volker Tresp
We address the learning of trust based on past observations and context information. We argue that from the truster's point of view trust is best expressed as one of several relations that exist between the agent to be trusted (trustee) and the state of the environment. Besides attributes expressing trustworthiness, additional relations might describe commitments made by the trustee with regard to the current situation, like: a seller offers a certain price for a specific product. We show how to implement and learn context-sensitive trust using statistical relational learning in form of the Infinite Hidden Relational Trust Model (IHRTM). The practicability and effectiveness of our approach is evaluated empirically on user-ratings gathered from eBay. Our results suggest that (i) the inherent clustering achieved in the algorithm allows the truster to characterize the structure of a trust-situation and provides meaningful trust assessments; (ii) utilizing the collaborative filtering effect associated with relational data does improve trust assessment performance; (iii) by learning faster and transferring knowledge more effectively we improve cold start performance and can cope better with dynamic behavior in open multiagent systems. The later is demonstrated with interactions recorded from a strategic two-player negotiation scenario.

128B
The Conclusion of Contracts by Software Agents in the Eyes of the Law
Tina Balke, Torsten Eymann
Faced with the ongoing evolution of software agents from mere passive tools to e-tailers acting autonomously for their human owners (principals), new legal challenges appear on the agenda. One of them is the question, whether the traditional law of agency that regulates the legal issues arising from human agents constituting legal relations between their principal and a third party, is applicable for software agents as well.
Based on the characteristics of software agents this paper examines approaches for a legal classification of software agents and thereby analyzes current legislation that deal with the conclusion of contracts by software agents.
Finally, this paper addresses remaining legal questions and discusses proposed solutions.

131B
Norm Emergence Under Constrained Interactions in Diverse Societies
Partha Mukherjee, Stephane Airiau, Sandip Sen
Effective norms, emerging from sustained individual interactions over time, can complement societal rules and significantly enhance performance of individual agents and agent societies. Researchers have used a model that supports the emergence of social norms via learning from interaction experiences where each interaction is viewed as a stage game. In this social learning model, which is distinct from an agent learning from repeated interactions against the same player, an agent learns a policy to play the game from repeated interactions with multiple learning agents. The key research question is to characterize when and how the entire population of homogeneous learners converge to a consistent norm when multiple action combinations yield the same optimal payoff. In this paper we study two extensions to the social learning model that significantly enhances its applicability. We first explore the effects of heterogeneous populations where different agents may be using different learning algorithms. We also investigate norm emergence when agent interactions are physically constrained. We consider agents located on a grid where an agent is more likely to interact with other agents situated closer to it than those that are situated afar. The key new results include the surprising acceleration in learning with limited interaction ranges. We also study the effects of pure-strategy players, i.e., nonlearners in the environment.

134B
Checking Correctness of Business Contracts via Commitments
Nirmit Desai, Nanjangud Narendra, Munindar Singh
Business contracts tend to be complex. In current practice, contracts are often designed by hand and

adopted by their participants after, at best, a manual analysis. This paper motivates and formalizes two aspects of contract correctness from the perspective of the preferences of the agents participating in them. A contract is safe for a participant if participating in the contract would not leave the participant worse off than otherwise. More strongly, a contract is beneficial to a participant if participating in the contract would leave the participant better off than otherwise.
This paper seeks to partially automate reasoning about the correctness of formally modeled business contracts. It represents contracts formally as a set of commitments. It motivates constraints on how cooperative agents might value the various states of commitments. Further, it shows that such constraints are consistent and promote cooperation. Lastly, it presents algorithms for checking the safety and guaranteed benefits of a contract.

Economic Paradigm I

Chair: Han La Poutre

220A
Strategic Betting for Competitive Agents
Liad Wagman, Vincent Conitzer
In many multiagent settings, each agent's goal is to come out ahead of the other agents on some metric, such as the currency obtained by the agent. In such settings, it is not appropriate for an agent to try to maximize its expected score on the metric; rather, the agent should maximize its expected probability of winning. In principle, given this objective, the game can be solved using game-theoretic techniques. However, most games of interest are far too large and complex to solve exactly. To get some intuition as to what an optimal strategy in such games should look like, we introduce a simplified game that captures some of their key aspects, and solve it (and several variants) exactly.
Specifically, the basic game that we study is the following: each agent i chooses a lottery over nonnegative numbers whose expectation is equal to her budget b_i. The agent with the highest realized outcome wins (and agents only care about winning). We show that there is a unique symmetric equilibrium when budgets are equal. We proceed to study and solve extensions, including settings where agents must obtain a minimum outcome to win; where agents choose their budgets (at a cost); and where budgets are private information.

228C
Synthesis of Strategies from Interaction Traces
Tsz-Chiu Au, Dana Nau, Sarit Kraus
We describe how to take a set of interaction traces produced by different pairs of players in a two-player repeated game, and combine them into a composite strategy. We provide an algorithm that, in polynomial time, can generate the best such composite strategy. We describe how to incorporate the composite strategy into an existing agent, as an enhancement of the agent's original strategy.
We provide experimental results using interaction traces from 126 agents (most of them written by students as class projects) for the Iterated Prisoner's Dilemma, Iterated Chicken Game, and Iterated Battle of the Sexes. We compared each agent with the enhanced version of that agent produced by our algorithm. The enhancements improved the agents' scores by about 5% in the IPD, 11% in the ICG, and 26% in the IBS, and improved their rank by about 12% in the IPD, 38% in the ICG, and 33% in the IBS.

236B
Artificial agents learning human fairness
Steven De Jong, Karl Tuyls, Katja Verbeeck
Recent advances in technology allow multi-agent systems to be deployed in cooperation with or as a service for humans. Example domains include resource distribution facilities, bargaining, aircraft deicing, etc. Typically, those systems are designed assuming perfectly rational, self-interested agents, according to the principles of classical game theory. However, research in the field of behavioral economics has shown that humans are not purely self-interested: they strongly care about fairness. Therefore, multi-agent systems that fail to take fairness into account, may not be sufficiently aligned with human expectations and may not reach intended goals. In this paper, we present a computational model for achieving fairness in adaptive multi-agent systems. The model uses a combination of Continuous Action Learning Automata and the Homo Egualis utility function. The novel contribution of our work is that this function is used in an explicit, computational manner. We show that results obtained by our 'fair' agents are compatible with experimental and analytical results obtained in the field of behavioral economics.

243C
Self-Interested Database Managers Playing The View Maintenance Game
Hala Mostafa, Victor Lesser, Gerome Miklau
A database view is a dynamic virtual table composed of the result set of a query, often executed over different underlying databases. The view maintenance problem concerns how a view is refreshed when the data sources are updated. We study the view maintenance problem when self-interested database managers from different institutions are involved, each concerned about the privacy of its database. We regard view maintenance as an incremental, sequential process where an action taken at a stage affects what happens at later stages. The contribution of this paper is twofold. First, we formulate the view maintenance problem as a sequential game of incomplete information where at every stage, each database manager decides what information to disclose, if any, without knowledge of the number or nature of updates at other managers. This allows us to adopt a satisficing approach where the final view need not reflect 100% of the databases updates. Second, we present an anytime algorithm for calculating epsilon Bayes-Nash equilibria that allows us to solve the large games which our problem translates to. Our algorithm is not restricted to games originating from the view maintenance problem; it can be used to solve general games of incomplete information. In addition, experimental results demonstrate our algorithm's attractive anytime behavior, which allows it to find good-enough solutions to large games within reasonable amounts of time.

248B
Zero-Intelligence Agents in Prediction Markets
Abraham Othman
We construct a novel agent-based model of prediction markets in which putative human qualities like learning, reasoning, and profit-seeking are absent. We show that the prices which emerge from a market populated by a class of distinctly inhuman agents, Zero-Intelligence agents with diffuse beliefs, replicate the findings of

empirical market studies. We use this result to argue against the prevailing descriptive theories of price formation in prediction markets, which have stressed the role of expert, rational participants.

252C
The Effects of Market-Making on Price Dynamics
Sanmay Das
This paper studies market-makers, agents responsible for maintaining liquidity and orderly price transitions in markets. Market-makers include major firms making markets on global stock exchanges, as well as software agents that run behind the scenes on novel electronic markets like prediction markets. We use a sophisticated model of market-making to build richer agent-based models of markets and show how these models can be useful both in understanding properties of existing markets and in predicting the impacts of structural changes. For example, we show how competition among market-makers can lead to significantly faster price discovery following a jump in the true value of an asset. We also show that myopic profit-maximization, apart from leading to poor market quality, is sub-optimal even for a monopolistic market-maker. This observation leads to an interesting characterization of the market-maker's exploration-exploitation dilemma as a tradeoff between price discovery and profit-taking.

Wednesday, May 14th, 13:40-15:00 – Session 2

Agent Theories, Models and Architectures I

Chair: Onn Shehory

97A
Quantifying Over Coalitions in Epistemic Logic
Thomas Agotnes, Wiebe van der Hoek, Michael Wooldridge
Some natural epistemic properties which may rise in applications can in standard epistemic logic only be expressed by formulas which are exponentially long in the number of agents in the system. For example, the expression of "at least m agents know that at most n agents know A" in standard epistemic logic is exponentially long in both n and the total number of agents in the system. We present Epistemic Logic with Quantification over Coalition (ELQC), where the standard common knowledge operator has been replaced giving expressions of the form $\langle P \rangle_{C\varphi}$ and $[P]_{C\varphi}$ where *P* is a *coalition predicate*, meaning that there is a coalition satisfying *P* which have common knowledge of φ and that all coalitions satisfying *P* has common knowledge of φ, respectively, and similarly for distributed knowledge and everybody-knows. The language is equally expressive but exponentially more succinct compared to standard epistemic logic. We give a sound and complete axiomatisation and a complete characterisation of the complexity of the model checking problem.

100A
Coalitions and Announcements
Thomas Agotnes, Hans van Ditmarsch
Two currently active strands of research on logics for multi-agent systems are dynamic epistemic logic, focusing on the epistemic consequences of actions, and logics of coalitional ability, focusing on what coalitions of agents can achieve by cooperating strategically. In this paper we make a first attempt to bridge these topics by considering the question: "what can a coalition achieve by public announcements?". We propose, first, an extension of public announcement logic with constructs of the form $\langle G \rangle \varphi$, where G is a set of agents, with the intuitive meaning that G can jointly make an announcement such that phi will be true afterwards. Second, we consider a setting where all agents can make (truthful) announcements at the same time, and propose a logic with a construct $\langle [G] \rangle \varphi$, meaning that G can jointly make an announcement such that no matter what the other agents announce, phi will be true. The latter logic is closely related to Marc Pauly's Coalition Logic.

103A
Modelling Coalitions: ATL + Argumentation
Nils Bulling, Carlos I. Chesnevar, Jurgen Dix
In the last few years, argumentation frameworks have been successfully applied to multi agent systems. Recently, argumentation has been used to provide a framework for reasoning about coalition formation.
At the same time alternating-time temporal logic has been used to reason about the behavior and abilities of coalitions of agents. However, ATL operators account only for the existence of successful strategies of coalitions. They do not consider whether coalitions can be actually formed.
This paper is an attempt to combine both frameworks and to develop a logic, where we can reason at the same time (1) about abilities of coalition of agents and (2) about the formation of coalitions. We provide a formal extension of ATL, in which the actual computation of the coalition is modelled in terms of argumentation semantics, which are expressed as first-class citizens within the logic. We show that its proof theory can be understood as a natural extension of the model checking procedure used in ATL.

106A
Information-based Deliberation
Carles Sierra, John Debenham
Information-based agency is founded on two observations: everything in an agent's world model is uncertain, and everything that an agent communicates gives away valuable information. The agent's deliberative mechanism manages interaction using plans and strategies in the context of the relationships the agent has with other agents, and is the means by which those relationships develop.

Agent Cooperation I

Chair: Eugénio Oliveira

162C
Regulating Air Traffic Flow with Coupled Agents
Adrian Agogino, Kagan Tumer
The ability to provide flexible, automated management of air traffic is critical to meeting the ever increasing needs of the next generation air transportation systems. This problem is particularly complex as it requires the integration of many factors including updated information (e.g., changing weather data), conflicting priorities (e.g., different airlines), limited resources (e.g., air traffic controllers) and very heavy traffic volume (e.g., over 40,000 flights over the US airspace). Furthermore, because the FAA will not accept black-box solutions, algorithmic improvements need to be consistent with current operating practices and provide explanations for each new decision. Unfortunately current methods

provide neither flexibility for future upgrades, nor high enough performance in complex coupled air flow management problems.
This paper extends agent-based methods for controlling air traffic flow to more realistic domains that have coupled flow patterns and need to be controlled through a variety of mechanisms. First, we explore an agent control structure that allows agents to control air traffic flow through one of three mechanisms (miles in trail, ground delays and rerouting). Second, we explore a new agent learning algorithm that can efficiently handle coupled flow patterns. We then test this agent solution on a series of congestion problems, showing that it is flexible enough to achieve high performance with different control mechanisms. In addition the results show that the new solution is able to achieve up to a 20% increase in performance over previous methods that did not account for the agent coupling.

165C
Decentralized Algorithms for Collision Avoidance in Airspace
David Sislak, Jiri Samek, Michal Pechoucek
The paper proposes decentralized deconfliction algorithms deployed on multiple autonomous aerial vehicles participating in free-flight operations. The paper provides two separate algorithms to collision avoidance - one based on iterative peer-to-peer negotiation solving a singular collision and second based on multi-party negotiation about a cluster of collisions. The presented decentralized algorithms allow the vehicles operating in the same area to utilize the given airspace more efficiently. The algorithms have been developed and tested on a multi-agent prototype and the properties of both algorithms are discussed on a set of large scale experiments.

168C
Heuristics for Negotiation Schedules in Multi-plan Optimization
Bo An, Fred Douglis, Fan Ye
In cooperating systems such as grids [4] and collaborative streaming Analysis [2], autonomous sites can establish “agreements” to arrange access to remote resources for a period of time [1]. The determination of which resources to reserve to accomplish a task need not be known a priori, because there exist multiple plans for accomplishing the same task and they may require access to different resources [3]. While these plans can be functionally equivalent, they may have different performance/cost tradeoffs and may use a variety of resources, both local and belonging to other sites. The negotiation schedule, i.e., the order in which remote resources are egotiated, determines how quickly one plan can be selected and deployed; it also decides the utility for running the plan. This paper studies the problem of optimizing negotiation schedules in cooperative systems with multiple alternative plans. We first provide a voting-based heuristic that reduces the complexity O(n!) of the exhaustive search to O(mnq). We also present a weight-based heuristic that further reduces the complexity to O(mn). Experimental results show that, on average, 1) the voting-based approach achieved 6% higher utility than the weight-based approach but the voting-based approach has much higher computation cost than the weight-based approach, 2) the two proposed approaches achieved almost 50% higher utility than a randomized approach; and 3) the average utility produced by the two proposed approaches are within almost 90% of that of the optimal results with reasonable plan sizes.

171C
Reaction Functions for Task Allocation to Cooperative Agents
Xiaoming Zheng, Sven Koenig
In this paper, we present our initial effort at solving task-allocation problems where cooperative agents need to perform tasks simultaneously. An example is multi-agent routing tasks where several agents need to visit targets simultaneously, for example, to move obstacles out of the way cooperatively. First, we propose reaction functions as a novel way of characterizing the costs of single agents in a distributed way. Second, we show how to approximate reaction functions so that their computation times are polynomial and their communication times are constant. Third, we show how reaction functions can be used by a central planner to allocate tasks to agents and determine their execution times. Finally, we show experimentally that the resulting task allocations are better than those of other greedy methods that do not use reaction functions.

Agent and Multi-Agent Learning I

Chair: Han La Poutre

194B
Autonomous Transfer for Reinforcement Learning
Matthew Taylor, Gregory Kuhlmann, Peter Stone
Recent work in transfer learning has succeeded in making reinforcement learning algorithms more efficient by incorporating knowledge from previous tasks. However, such methods typically must be provided either a full model of the tasks or an explicit relation mapping one task into the other. An autonomous agent may not have access to such high-level information, but would be able to analyze its experience to find similarities between tasks. In this paper we introduce Modeling Approximate State Transitions by Exploiting Regression (MASTER), a method for automatically learning a mapping from one task to another through an agent's experience. We empirically demonstrate that such learned relationships can significantly improve the speed of a reinforcement learning algorithm in a series of Mountain Car tasks. Additionally, we demonstrate that our method may also assist with the difficult problem of task selection for transfer.

197B
Analysis of an Evolutionary Reinforcement Learning Method in a Multiagent Domain
Jan Hendrik Metzen, Mark Edgington, Yohannes Kassahun, Frank Kirchner
Many multiagent problems comprise subtasks which can be considered as reinforcement learning (RL) problems. In addition to classical temporal difference methods, evolutionary algorithms are among the most promising approaches for such RL problems. The relative performance of these approaches in certain subdomains (e.g. multiagent learning) of the general RL problem remains an open question at this time. In addition to theoretical analysis, benchmarks are one of the most important tools for comparing different RL methods in certain problem domains. A recently proposed multiagent RL benchmark problem is the RoboCup Keepaway benchmark. This benchmark is one of the most challenging multiagent learning problems because its

state-space is continuous and high dimensional, and both the sensors and the actuators are noisy. In this paper we analyze the performance of the neuroevolutionary approach called Evolutionary Acquisition of Neural Topologies (EANT) in the Keepaway benchmark, and compare the results obtained using EANT with the results of other algorithms tested on the same benchmark.

200B
The Utility of Temporal Abstraction in Reinforcement Learning
Nicholas Jong, Todd Hester, Peter Stone
The hierarchical structure of real-world problems has motivated extensive research into temporal abstractions for reinforcement learning, but precisely how these abstractions allow agents to improve their learning performance is not well understood. This paper investigates the connection between temporal abstraction and an agent's exploration policy, which determine how the agent's performance improves over time. Experimental results with standard methods for incorporating temporal abstractions show that these methods benefit learning only in limited contexts. The primary contribution of this paper is a clearer understanding how hierarchical decompositions interact with reinforcement learning algorithms, with important consequences for the manual design or automatic discovery of action hierarchies.

203B
Switching Dynamics of Multi-Agent Learning
Peter Vrancx, Karl Tuyls, Ronald Westra
This paper presents the dynamics of multi-agent reinforcement learning in multiple state problems. We extend previous work that formally modelled the relation between reinforcement learning agents and replicator dynamics in stateless multi-agent games. More precisely, in this work we use a combination of replicator dynamics and switching dynamics to model multi-agent learning automata in multi-state games. This is the first time that the dynamics of problems with more than one state is considered with replicator equations. Previously, it was unclear how the replicator dynamics of stateless games had to be extended to account for multiple states. We use our model to visualize the basin of attraction of the learning agents and the boundaries of switching dynamics at which an agent possibly arrives in a new dynamical system. Our model allows to analyze and predict the behavior of the different learning agents in a wide variety of multi-state problems. In our experiments we illustrate this powerful method in two games with two agents and two states.

Agent Reasoning I

Chair: Rafael Bordini

141C
Goal Generation with Relevant and Trusted Beliefs
Célia da Costa Pereira, Andrea G. B. Tettamanzi
A rational agent adopts (or changes) its *goals* when new information (*beliefs*) becomes available or its *desires* (e.g., tasks it is supposed to carry out) change. In conventional approaches to goal generation in which a goal is considered as a "particular" desire, a goal is adopted if and only if *all* conditions leading to its generation are satisfied. It is then supposed that all beliefs are both equally *relevant* and completely *trusted*. However, that is not a realistic setting. In fact, depending on the agent's trust in the source of a piece of information, an agent may decide how strongly it takes into consideration such piece of information in goal generation. On the other hand, not all beliefs are equally relevant to the adoption of a given goal, and a given belief may not be equally relevant to the adoption of different goals.
We propose an approach which takes into account both the relevance of beliefs and the trust degree of the source from which the corresponding piece of information comes, in desire/goal generation.

144C
Suspending and Resuming Tasks in Intelligent Agents
John Thangarajah, James Harland, David Morley, Neil Yorke-Smith
Intelligent agents designed to work in complex, dynamic environments must respond robustly and flexibly to environmental and circumstantial changes. An agent must be capable of deliberating about appropriate courses of action, which may include re-prioritising goals, aborting particular tasks, or scheduling tasks in a particular order. This paper investigates the incorporation of a mechanism to suspend and reconsider tasks within a BDI-style architecture. Such an ability provides an agent designer greater flexibility to direct agent operation, and it offers a generic means for handling conflicts between tasks. We investigate conditions under which a goal or a plan may be suspended, the process for suspending it, and the appropriate behaviours upon resumption. We give an operational semantics for suspending tasks in terms of the abstract agent language CAN, thus providing a general mechanism that can be incorporated into any BDI-based agent programming language.

147C
Partial Goal Satisfaction and Goal Change
Yi Zhou, Leon van der Torre, Yan Zhang
Partial implication semantics in the context of a background theory has been introduced to formalize partial goal satisfaction in the context of beliefs. In this paper, we introduce strong partial implication prohibiting *redundancies* and weak partial implication allowing *side effects*, we study their semantic as well as complexity properties, and we apply the three notions of partial implication to goal change in the context of beliefs.

150C
Belief Operations for Motivated BDI Agents
Patrick Krümpelmann, Matthias Thimm, Gabriele Kern-Isberner, Manuela Ritterskamp
The beliefs of an agent reflecting her subjective view of the world constitute one of the main components of a BDI agent. In order to incorporate new information coming from other agents, or to adjust to changes in the environment, the agent has to carry out belief change operations while taking meta-logical information on time and reliabilities into account. In this paper, describe a framework for belief operations within a BDI agent, sketching the interactions of beliefs with desires and intentions, respectively. Furthermore, we illustrate how motivations and know-how come into play in our agent model of this framework. We focus on the presentation of a complex setting for belief change that makes use of techniques both from merging and update, and provides a BDI agent with advanced reasoning capabilities. Extended logic programs under the answer set semantics

will serve as the basic knowledge representation formalism.

Economic Paradigms II

Chair: Kate Larson

221B

Playing Games for Security: An Efficient Exact Algorithm for Solving Bayesian Stackelberg Games

Praveen Paruchuri, Jonathan Pearce, Janusz Marecki, Milind Tambe, Fernando Ordonez, Sarit Kraus

In a class of games known as Stackelberg games, one agent (the leader) must commit to a strategy that can be observed by the other agent (the follower or adversary) before the adversary chooses its own strategy. We consider Bayesian Stackelberg games, in which the leader is uncertain about the types of adversary it may face. Such games are important in security domains, where, for example, a security agent (leader) must commit to a strategy of patrolling certain areas, and a robber (follower) has a chance to observe this strategy over time before choosing its own strategy of where to attack. This paper presents an efficient exact algorithm for finding the optimal strategy for the leader to commit to in these games. This algorithm, DOBSS, is based on a novel and compact mixed-integer linear programming formulation. Compared to the most efficient algorithm known previously for this problem, DOBSS is not only faster, but also leads to higher quality solutions, and does not suffer from problems of infeasibility that were faced by this previous algorithm. Note that DOBSS is at the heart of the ARMOR system that is currently being tested for security scheduling at a major international airport.

229A

Solving two-person zero-sum repeated games of incomplete information

Andrew Gilpin, Tuomas Sandholm

In repeated games with incomplete information, rational agents must carefully weigh the tradeoffs of advantageously exploiting their information to achieve a short-term gain versus carefully concealing their information so as not to give up a long-term informed advantage. The theory of infinitely-repeated two-player zero-sum games with incomplete information has been carefully studied, beginning with the seminal work of Aumann and Maschler. While this theoretical work has produced a characterization of optimal strategies, algorithms for solving for optimal strategies have not yet been studied. For the case where one player is informed about the true state of the world and the other player is uninformed, we provide a nonconvex mathematical programming formulation for computing the value of the game, as well as optimal strategies for the informed player. We then describe an efficient algorithm for solving this difficult optimization problem to within arbitrary accuracy. We also describe an efficient algorithm for finding optimal strategies for the uninformed player.

237C

A heads-up no-limit Texas Hold'em poker player: Discretized betting models and automatically generated equilibrium finding programms

Andrew Gilpin, Tuomas Sandholm, Troels Bjerre Sørensen

We present a game theory-based player for heads-up no-limit Texas Hold'em poker. Our player is built from three components. First, to deal with the virtually infinite strategy space of no-limit poker, we develop a discretized betting model designed to capture the most important strategic choices in the game. Second, we employ potential-aware automated abstraction algorithms for identifying strategically similar situations in order to decrease the size of the game tree. Third, we develop a new technique for automatically generating the source code of an equilibrium-finding algorithm from an XML-based description of a game. This automatically generated program is more efficient than what would be possible with a general-purpose equilibrium-finding program. Finally, we present experiments demonstrating that our player is competitive with the best existing computer opponents.

****244A****

Computing an Approximate Jam/Fold Equilibrium for 3-player No-Limit Texas Hold'em Tournaments

Sam Ganzfried, Tuomas Sandholm

A recent paper computes near-optimal strategies for two-player no-limit Texas hold'em tournaments; however, the techniques used are unable to compute equilibrium strategies for tournaments with more than two players. Motivated by the widespread popularity of multiplayer tournaments and the observation that jam/fold strategies are near-optimal in the two player case, we develop an algorithm that computes approximate jam/fold equilibrium strategies in tournaments with three – and potentially even more – players. Our algorithm combines an extension of fictitious play to imperfect information games, an algorithm similar to value iteration for solving stochastic games, and a heuristic from the poker community known as the Independent Chip Model which we use as an initialization. Several ways of exploiting suit symmetries and the use of custom indexing schemes made the approach computationally feasible. Aside from the initialization and the restriction to jam/fold strategies, our high level algorithm makes no poker-specific assumptions and thus also applies to other multiplayer stochastic games of imperfect information.

Wednesday, May 14th, 15:30-16:30 – Session 3

IFAAMAS-07 Victor Lesser Distinguished Dissertation Award

Truthful Reputation Mechanisms for Online Systems

Radu Jurca

Multi-Robotics I

Chair: Pedro Lima

37A

Adaptive Multi-Robot Wide-Area Exploration and Mapping

Kian Hsiang Low, John Dolan, Pradeep Khosla

The exploration problem is a central issue in mobile robotics. A complete terrain coverage is not practical if the environment is large with only a few small hotspots. This paper presents an adaptive multi-robot exploration strategy that is novel in performing both wide-area coverage and hotspot sampling using non-myopic path planning. As a result, the environmental phenomena can be accurately mapped. It is based on a dynamic programming formulation, which we call the Multi-robot

Adaptive Sampling Problem (MASP). A key feature of MASP is in covering the entire adaptivity spectrum, thus allowing strategies of varying adaptivity to be formed and theoretically analyzed in their performance; a more adaptive strategy improves mapping accuracy. We apply MASP to sampling the Gaussian and log-Gaussian processes, and analyze if the resulting strategies are adaptive and maximize wide-area coverage and hotspot sampling. Solving MASP is non-trivial as it comprises continuous state components. So, it is reformulated for convex analysis, which allows discrete-state monotone-bounding approximation to be developed. We provide a theoretical guarantee on the policy quality of the approximate MASP (aMASP) for using in MASP. Although aMASP can be solved exactly, its state size grows exponentially with the number of stages. To alleviate this computational difficulty, anytime algorithms are proposed based on aMASP, one of which can guarantee its policy quality for MASP in real time.

40A

A Decentralized Approach to Cooperative Situation Assessment in Multi-Robot Systems

Giuseppe Settembre, Alessandro Farinelli, Paul Scerri, Katia Sycara, Daniele Nardi

To act effectively under uncertainty, multi-robot teams need to accurately estimate the state of the environment. Although individual robots, with uncertain sensors, may not be able to accurately determine the current situation, the team as a whole should have the capability to perform situation assessment. However, to share all information acquired by each single robot is not an acceptable solution in many scenarios. This paper presents a decentralized approach to cooperative situation assessment. In particular, robots have to cooperatively decide which team plan to execute given their information. Plan selection is performed using a decision theoretical approach. When a robot selects a plan, it sends a proposal for that plan, to one of its team mates. A robot receiving a plan proposal, can either agree with the plan forwarding the proposal on, or it can provide sensor information to suggest that an alternative plan might have higher expected utility. Once sufficient robots agree with the proposal, the plan is instantiated. The algorithm successfully balances the value of cooperative sensing against the cost of sharing large volumes of information. Experiments verify the utility of the approach , showing that the algorithm dramatically out-performs individual decision-making and obtains performance similar to a centralized approach.

43A

Self-organized Flocking with a Mobile Robot Swarm

Ali Turgut, Hande Celikkanat, Fatih Gokce, Erol Sahin

This paper studies self-organized flocking in a swarm of mobile robots. We present Kobot, a mobile robot platform developed specifically for swarm robotic studies briefly describing its sensing and communication abilities. In particular, we describe a scalable method that allows the robots to sense the orientations of their neighbors using a digital compass and wireless communication. Then we propose a behavior for a swarm of Kobots that creates self-organized flocking by using heading alignment and proximal control. The flocking behavior is observed to operate in three phases: alignment, advance, and avoidance. We evaluate four variants of this behavior by setting its parameters to extremum values and analyze the performance of flocking using a number of metrics, such as order and entropy. Our results show that, the flocking behavior obtained under appropriate parameter values, is quite robust and generates successful self-organized flocking in environments with obstacles.

Agent Reasoning II

Chair: Mehdi Dastani

153C

A constrained argumentation system for practical reasoning

Leila Amgoud, Caroline Devred, Marie-Christine Lagasquie-Schiex

Practical reasoning (PR), which is concerned with the generic question of what to do, is generally seen as a two steps process: (1) deliberation, in which an agent decides what state of affairs it wants to reach – that is, its desires; and (2) means-ends reasoning, in which the agent looks for plans for achieving these desires. A desire is justified if it holds in the current state of the world, and feasible if there is a plan for achieving it. The agent's intentions are thus a consistent subset of desires that are both justified and feasible. This paper proposes the first argumentation system for PR that computes in one step the intentions of an agent, allowing thus to avoid the drawbacks of the existing systems. The proposed system is grounded on a recent work on constrained argumentation systems, and satisfies the rationality postulates identified in argumentation literature, namely the *consistency* and the *completeness* of the results.

156C

Using Enthymemes in an Inquiry Dialogue System

Elizabeth Black, Anthony Hunter

A common assumption for logic-based argumentation is that an argument is a pair <Phi, alpha> where Phi is a minimal subset of the knowledgebase such that Phi is consistent and Phi entails the claim alpha. However, real arguments (i.e. arguments presented by humans) usually do not have enough explicitly presented premises for the entailment of the claim (i.e. they are enthymemes). This is because there is some common knowledge that can be assumed by a proponent of an argument and the recipient of it. This allows the proponent of an argument to encode an argument into a real argument by ignoring the common knowledge, and it allows a recipient of a real argument to decode it into the intended argument by drawing on the common knowledge. If both the proponent and recipient use the same common knowledge, then this process is straightforward. Unfortunately, this is not always the case, and this raises interesting issues for dialogue systems in which the recipient has to cope with the disparities between the different views on what constitutes common knowledge. Here we investigate the use of enthymemes in inquiry dialogues. For this, we propose a generative inquiry dialogue system and show how, in this dialogue system, enthymemes can be managed by the agents involved, and how common knowledge can evolve through dialogue.

142A

A dialogue mechanism for public argumentation using conversation policies

Yuqing Tang, Simon Parsons

Most existing dialogue systems that are based on argumentation connect the state of an agents' beliefs to the conversations that an agent may carry out. In this

paper we propose an alternative, decoupling the mechanism that an agent uses to maintain the coherence of its beliefs from the mechanism that the agent uses to carry out a dialogue. The advantage of such an arrangement is that is provides greater flexibility in designing agents. Under our proposal, a given agent can use a range of different conversation policies, some of which require coherence between its utterances and its beliefs, and some of which do not. This, in turn, makes it easier to create agents that operate in a range of domains and thus need to undertake a range of different kinds of dialogue, dialogues that make use of a range of different kinds of domain knowledge.

Agent Cooperation II

Chair: Katia Sycara

174C
Role-Based Teamwork Activity Recognition in Observations of Embodied Agent Actions
Linus Luotsinen, Ladislau Boloni
Recognizing team actions in the behavior of embodied agents has many practical applications and had seen significant progress in recent years. One approach with proven results is based on HMM-based recognition of spatio-temporal patterns in the behavior of the agents. While it had been shown to work in real-world datasets, this approach was found to be brittle.
In this paper we present two contributions which together can significantly increase the robustness of teamwork activity recognition. First we introduce a technique to reduce high dimensional continuous input data to a set of discrete features, which capture the essential components of the team actions. Second, we prefix the actual team action recognition with a role recognition module, which allows us to present the recognizer with arbitrarily shuffled input, and still obtain high recognition rates.
We validate the improved accuracy and robustness of the team action recognizer on datasets derived from captured real world data.

163A
Tags and Image Scoring for Robust Cooperation
Nathan Griffiths
Establishing and maintaining cooperation is an enduring problem in multi-agent systems and, although several solutions exist, the increased use of online trading systems, peer-to-peer networks, and ubiquitous computing environments mean that it remains an important question. Environments are emerging in which large numbers of agents are required to cooperate, but where repeat interactions between agents may be rare or non-existent. Most existing approaches to cooperation rely on reciprocity to establish notions of trust and reputation. However, where repeat interactions are rare such approaches are not always effective. In this paper we use ideas from biology and the social sciences to provide a mechanism that supports cooperation in such environments. Our mechanism combines a tag-based method to enable cooperation given a lack of reciprocity, with an adaptation of a simple image scoring reputation model to cope with cheating agents. Using a simple peer-to-peer scenario we show how cooperative behaviour is favoured, and how the influence of cheating agents can be reduced using only minimal information about an agent's neighbours.

166A
Efficiently Determining the Appropriate Mix of Personal Interaction and Reputation Information in Partner Choice
Shulamit Reches, Philip Hendrix, Barbara Grosz, Sarit Kraus
Many multi-agent settings require that agents identify appropriate partners or teammates with whom to work on tasks. In selecting among potential partners, agents may benefit from obtaining information about the alternative possibilities through gossip (i.e., by consulting others) or using a reputation system (a centralized repository of information about past behavior). This paper defines a statistical model, the "Information-Acquisition Source Utility (IASU) model" by which agents operating in an uncertain world can determine the amount of information to collect about potential partners before choosing one and which information sources they should consult (gossip, reputation system, or additional personal interaction with the agent). The IASU model represents the cost of information, which may vary by information source. To maximize the expected gain from the choice, it estimates the utility of choosing a partner by iteratively estimating the benefit of additional information. The paper reports empirical studies that compare the effectiveness of the IASU model with a baseline in which only prior experience with a potential partner is used as the basis of the decision and with a model that determines in advance both the amount of information and the allocation among different resources. Two different application domains are used in these empirical studies, the Surrogate Venture Game model, which deals with choosing an optimal partner for a business venture, and a restaurant domain. The results of the experiments show that the use of the model significantly increases agents' overall utility.

Economic Paradigm III

Chair: Jeff Rosenschein

****222C****
Anonymity-Proof Shapley Value: Extending Shapley Value for Coalitional Games in Open Environments
Naoki Ohta, Vincent Conitzer, Yasufumi Satoh, Atsushi Iwasaki, Makoto Yokoo
Coalition formation is an important capability for automated negotiation among self-interested agents. In order for coalitions to be stable, a key question that must be answered is how the gains from cooperation are to be distributed. Coalitional game theory provides a number of solution concepts for this. However, recent research has revealed that these traditional solution concepts are vulnerable to various manipulations in open anonymous environments such as the Internet. To address this, previous work has developed a solution concept called the anonymity-proof core, which is robust against such manipulations. That work also developed a method for compactly representing the anonymity-proof core. However, the required computational and representational costs are still huge.
In this paper, we develop a new solution concept which we call the anonymity-proof Shapley value. We show that the anonymity-proof Shapley value is characterized by certain simple axiomatic conditions, always exists, and is uniquely determined. The computational and representational costs of the anonymity-proof Shapley

value are drastically smaller than those of existing anonymity-proof solution concepts.

230B
An anytime approximation method for the inverse Shapley value problem
Shaheen Fatima, Michael Wooldridge, Nick Jennings
Coalition formation is a key form of interaction in multi-agent systems. It is the process of bringing together two or more agents so as to achieve goals that individuals on their own cannot, or to achieve them more efficiently. Typically, in such situations, the agents have conflicting preferences over the set of possible joint goals. Thus, before the agents realize the benefits of cooperation, they must find a way of resolving these conflicts and reaching a consensus. In this context, cooperative game theory offers the voting game as a mechanism for agents to reach a consensus. It also offers ways of measuring the influence or power a player has in determining the outcome of a voting game. In such situations, the Shapley value is a way of measuring an agent's voting power. Given this, the designer of a voting game wants to construct a game such that a player's Shapley value is equal to some desired value. This is called the inverse Shapley value problem. Solving this problem is necessary, for instance, to ensure fairness in the players' voting powers. However, from a computational perspective, finding a player's Shapley value for a given game is #P-complete. Consequently, the problem of verifying that a voting game does indeed yield the required powers to the agents is also #P-complete. Therefore, in order to overcome this problem we present a computationally efficient approximation algorithm for solving the inverse problem. This method is based on the technique of ""successive approximations""; it starts with some initial approximate solution and iteratively updates it such that after each iteration, the approximate gets closer to the required solution. This is an anytime algorithm and has time complexity polynomial in the number of players. We also analyze the performance of this method in terms of its approximation error and the rate of convergence of an initial solution to the required one. Specifically, we show that the former decreases after each iteration, and that the latter increases with the number of players and also with the initial approximation error.

238A
Approximating Power Indices
Yoram Bachrach, Vangelis Markakis, Ariel Proccacia, Jeffrey Rosenschein, Amin Saberi
Many multiagent domains where cooperation among agents is crucial to achieving a common goal can be modeled as coalitional games. However, in many of these domains, agents are unequal in their power to affect the outcome of the game. Prior research on weighted voting games has suggested power indices, which reflect how much “real power” a voter has. Although primarily used for voting games, these indices can be applied to any simple coalitional game. Computing these indices is known to be computationally hard in various domains, so one must resort to approximate methods for calculating them.
We suggest and analyze randomized methods to approximate power indices such as the Banzhaf power index and the Shapley-Shubik power index. Our approximation algorithms do not depend on a specific representation of the game, so they can be used in *any* simple coalitional game. Our methods are based on testing the game's value for several sample coalitions. We also show that no approximation algorithm can do much better for general coalitional games, by providing lower bounds for both deterministic and randomized algorithms for calculating power indices.

Thursday, May 15th, 8:40-9:40 – Invited Talk

Game-Theoretic Recommendations: Some Progress in an Uphill Battle
Moshe Tennenholtz
In this talk we consider two highly challenging problems in the foundations of game theory and its application to multi-agent systems. Namely, we consider the question of how should an agent choose its action in a given game, and the task of leading agents to adopt desired behaviors in a given game. In the recent years we provided some useful attacks on these fundamental problems. Our study of competitive safety analysis and the study of learning in ensembles of games, provide surprisingly useful tools for addressing the first challenge. Our theory of mediators provides powerful tools for addressing the second challenge. These approaches refer to non-cooperative games; in the context of social choice, we introduced the axiomatic approach to ranking/reputation/trust/ recommendation systems; in particular, our work on trust-based recommendation systems introduces several basic results in the characterization of useful recommendation techniques.

Thursday, May 15th, 10:10-12:10 – Session 4

Agent Theories, Models and Architectures II

Chair: Rafael Bordini

98B
A Temporal Logic for Markov Chains
Wojtek Jamroga
Most models of agents and multi-agent systems include information about possible states of the system (that defines relations between states and their external characteristics), and information about relationships between states. Qualitative models of this kind assign no numerical measures to these relationships. At the same time, quantitative models assume that the relationships are measurable, and provide numerical information about the degrees of relations. In this paper, we explore the analogies between some qualitative and quantitative models of agents/processes, especially those between transition systems and Markovian models.
Typical analysis of Markovian models of processes refers only to the expected utility that can be obtained by the process. On the other hand, modal logic offers a systematic method of describing phenomena by combining various modal operators. Here, we try to exploit linguistic features, offered by propositional modal logic, for analysis of Markov chains and Markov decision processes. To this end, we propose Markov Temporal Logic - a multi-valued logic that extends the branching time logic CTL*.

101B
Model-Checking Agent Refinement
Lacramioara Astefanoaei, Frank S. de Boer
We present a proof-technique for reducing the nondeterminism of abstract agent specifications in a BDI framework by means of refinement. We implement the operational semantics of agent specifications in rewrite

systems, such that we can automatically check if refinement between (fair) executions of agents holds.

104B

Goals in Agent Systems: A Unifying Framework

M. Birna van Riemsdijk, Mehdi Dastani, Michael Winikoff

In the literature on agent systems, the proactive behavior of agents is often modeled in terms of goals that the agents pursue. We review a number of commonly-used existing goal types and propose a simple and general definition of goal, which unifies these goal types. We then give a formal and generic operationalization of goals by defining an abstract goal architecture, which describes the adoption, pursuit, and dropping of goals in a generic way. This operationalization is used to characterize the discussed goal types.

99C

Internal models and private multi agent belief revision

Guillaume Aucher

We generalize the AGM belief revision theory to the multi-agent case. To do so, we first generalize the semantics of the single-agent case, based on the notion of interpretation, to the multi-agent case. Then we show that, thanks to the shape of our new semantics, all the results of the AGM framework transfer. Afterwards we investigate some postulates that are specific to our multi-agent setting. Finally, we give an example of revision operator that fulfills one of these new postulates and give an example of revision on a concrete example.

102C

Anchoring Institutions in Agents' Attitudes: Towards a Logical Framework for Autonomous MAS

Benoit Gaudou, Dominique Longin, Emiliano Lorini, Luca Tummolini

The aim of this paper is to provide a logical framework for the specification of autonomous Multi-Agent Systems (MAS). A MAS is autonomous in so far as it is capable of binding (‘nomos’) itself (‘auto’) independently of any external normative constraint specified by a designer. In particular, a MAS is autonomous if it is able to maintain its social institutions (i.e. rule-governed social practices) only by way of the agents’ attitudes. In order to specify an autonomous MAS, we propose the logic AL (Acceptance Logic) in which the acceptance of a proposition by the agents qua group members (i.e. group acceptance) is introduced. Such propositions are true w.r.t. an institutional context and correspond to facts that are instituted in an attitude-dependent way (i.e. normative and institutional facts). Finally, we contend that the present approach paves the way for a foundation of legal institutions, for studying the interaction between social and legal institutions and, eventually, for understanding and modeling institutional change.

105C

Verifying time, memory and communication bounds in systems of reasoning agents

Natasha Alechina, Brian Logan, Hoang Nga Nguyen, Abdur Rakib

We present a framework for verifying systems composed of heterogeneous reasoning agents, in which each agent may have differing knowledge and inferential capabilities, and where the resources each agent is prepared to commit to a goal (time, memory and communication bandwidth) are bounded. The framework allows us to investigate, for example, whether a goal can be achieved if a particular agent, perhaps possessing key information or inferential capabilities, is unable (or unwilling) to contribute more than a given portion of its available computational resources or bandwidth to the problem. We present a novel temporal epistemic logic, BMCL, which allows us to describe a set of reasoning agents with bounds on time, memory and the number of messages they can exchange. The bounds on memory and communication are expressed as axioms in the logic. As an example, we show how to axiomatize a system of agents which reason using resolution and prove that the resulting logic is sound and complete. We then show how to encode a simple system of reasoning agents specified in BMCL in the description language of a model checker, and verify that the agents can achieve a goal only if they are prepared to commit certain time, memory and communication resources.

Agent Simulations / Emergent Behaviour

Chair: Elisabeth Sklar

110B

Resource Limitations, Transmission Costs and Critical Thresholds in Scale-Free Networks

Chung-Yuan Huang, Chuen-Tsai Sun, Chia-Ying Cheng, Yu-Shiuan Tsai

Whether or not a critical threshold exists when epidemic diseases are spread in complex networks is a problem attracting attention from researchers in several disciplines. In 2001, Pastor-Satorras and Vespignani used a computational simulations approach to show that epidemic diseases which spread through scale-free social networks do not have positive critical thresholds. In other words, even if a disease has almost no chance of being transmitted from one person to another, it can still spread throughout a scale-free network. However, they ignored two key factors that have a large impact on epidemic dynamics: economic resource limitations and transmission costs. Every infection event entails tangible or intangible costs in terms of time, energy, or money to the carrier, recipient, or both. Here we apply an agent-based modeling and network-oriented computer simulation approach to analyze the influences of resource limitations and transmission costs on epidemic dynamics and critical thresholds in scale-free networks. Our results indicate that when those resources and costs are taken into consideration, the epidemic dynamics of scale-free networks are very similar to those of homogeneous networks, including the presence of significant critical thresholds. It is hoped that our data will help epidemiologists, public health professionals, and computer scientists working with core questions of epidemic diseases, estimates of epidemic dynamics and spreading, and effective public health policies and immunization strategies.

113B

Agent-Based Simulation of the Spatial Dynamics of Crime

Tibor Bosse, Charlotte Gerritsen

An important challenge within the field of Criminology is to investigate the spatio-temporal dynamics of crime. Typical questions in this area are how the behaviour of offenders, targets, and guardians, and the emergence and displacement of criminal hot spots can be predicted. This paper presents an agent-based simulation model that can be used as an experimental tool to address such questions.

The simulation model particularly focuses on the interplay between hot spots and reputation. Using the model, a number of simulation experiments have been performed, of which results have been analysed using formal techniques. The results indicate that the presented approach is able to adequately reproduce displacement patterns as described in the literature.

116B

Reusing Models in Multi-Agent Simulation with Software Components

Paulo Salem da Silva, Ana Cristina Vieira de Melo

Simulation models are abstract representations of systems one wants to study through computer simulation. In multi-agent based simulation, such models usually represent agents and their relations. An important issue concerning these models is how they can be effectively reused across different simulations. But while much attention has been given to other engineering issues, model reuse has remained mostly untreated. To help address this issue, in this paper we present both a method and a software architecture for multi-agent simulation designed with reuse in mind. We employ software components as fundamental reusable model assets and show how their composition can also be reused. Our technique depends on some domain specific assumptions, such as the fact that agents must be related by social networks, and we argue that these are actually helpful in the context of software components. A case study is also given in order to illustrate clearly how the same component can be reused in two distinct simulation problems using our approach.

119B

Agent-based Models for Animal Cognition: A Proposal and Prototype

Elske van der Vaart, Rineke Verbrugge

Animal ecologists have successfully applied agent-based models to many different problems. Often, these focus on issues concerning collective behaviors, environmental interactions, or the evolution of traits. In these cases, patterns of interest can usually be investigated by constructing the appropriate multi-agent system, and then varying or evolving model parameters. In recent years, however, the study of animal behavior has increasingly expanded to include the study of animal cognition. In this field, the question is not just how or why a particular behavior is performed, but also what its 'mental underpinnings' are. In this paper, we argue that agent-based models are uniquely suited to explore questions concerning animal cognition, as the experimenter has direct access to agents' internal representations, control over their evolutionary history, and a perfect record of their previous learning experience. To make this possible, a new modeling paradigm must be developed, where agents' reasoning processes are explicitly simulated, and can evolve over time. We propose that this be done in the form of "if-then" rules, where only the form is specified, not the content. This should allow qualitatively different reasoning processes to emerge, which may be more or less "cognitive" in nature. In this paper, we illustrate the potential of such an approach with a prototype model. Agents must evolve explicit rule sets to forage for food, and to escape predators. It is shown that even in this relatively simple setup, different strategies emerge, as well as unexpected outcomes.

111C

Multi Agent Based Simulation of Transport Chains

Paul Davidsson, Johan Holmgren, Jan Persson, Linda Ramstedt

An agent-based tool for micro-level simulation of transport chains (TAPAS) is described. It is more powerful than traditional approaches as it is able to capture the interactions between individual actors of a transport chain, as well as their heterogeneity and decision making processes. Whereas traditional approaches rely on assumed statistical correlation between different parameters, TAPAS relies on causality, i.e., the decisions and negotiations that lead to the transports being performed. An additional advantage is that TAPAS is able to capture time aspects, such as, the influence of timetables, arrival times, and time-differentiated taxes and fees. TAPAS is composed of two layers, one layer simulating the physical activities taking place in the transport chain, e.g., production, storage, and transports of goods, and an-other layer simulating the different actors' decision making processes and interaction. The decision layer is implemented as a multi-agent system using the JADE platform, where each agent corresponds to a particular actor. We demonstrate the use of TAPAS by investigating how the actors in a transport chain are expected to act when different types of governmental control policies are applied, such as, fuel taxes, road tolls, and vehicle taxes. By analyzing the costs and environmental effects, TAPAS provides guidance in decision making regarding such control policies. We argue that TAPAS may also complement existing approaches in different ways, for instance by generating input data such as transport demand. Since TAPAS models a larger part of the supply chain, the transport demand is a natural part of the output. Studies may concern operational decisions like choice of consignment size and frequency of deliveries, as well as strategic decisions like where to locate storages, terminals, etc., choice of producer, and adaptation of vehicle fleets.

114C

Programming Agents as a Means of Capturing Self-Strategy

Michal Chalamish, David Sarne, Sarit Kraus

In this paper we report results of an extensive evaluation of people's ability to reproduce the strategies they use in simple real-life settings. Having the ability to reliably capture people's strategies in different environments is highly desirable in Multi-Agent Systems (MAS). However, as trivial and daily as these strategies are, the process is not straightforward and people often have a different belief of how they act. We describe our experiments in this area, based on the participation of a pool of subjects in four different games with variable complexity and characteristics. The main measure used for determining the closeness between the two types of strategies used is the level of similarity between the actions taken by the participants and those taken by agents they programmed in identical world states. Our results indicate that generally people have the ability to reproduce their game strategies for the class of games we consider. However, this process should be handled carefully as some individuals tend to exhibit a behavior different from the one they program into their agents. The paper evaluates one possible method for enhancing the process of strategy reproduction.

Industry and Applications Track I

Chair: Anna Perini

Autonomic Machine Control - A Case Study
Christian Dannegger, Dominic Greenwood
This paper describes an autonomic machine control system applied to the adaptive control of a modular soldering machine. The particular case concerns the creation of a novel modular production machine with an integrated distributed agent control system, which will be sold worldwide from mid 2008. The agent model is described in terms of the specific customer requirements and specific advantages of the approach discussed.

82A
An Interactive Platform for Auction-Based Allocation of Loads in Transportation Logistics
Valentin Robu, Han Noot, Han La Poutré, Willem-Jan van Schijndel
This paper describes an agent-based platform for the allocation of loads in distributed transportation logistics, developed as a collaboration between CWI, Dutch National Center for Mathematics and Computer Science, Amsterdam and Vos Logistics Organizing, Nijmegen, The Netherlands.
The platform follows a real business scenario proposed by Vos, and it involves a set of agents bidding for transportation loads to be distributed from a central depot in the Netherlands to different locations across Germany. The platform supports both human agents (i.e. transportation planners), who can bid through specialized planning and bidding interfaces, as well as automated, software agents. We exemplify how the proposed platform can be used to test both the bidding behaviour of human logistics planners, as well as the performance of automated auction bidding strategies, developed for such settings.
The paper first introduces the business problem setting and then describes the architecture and main characteristics of our auction platform. We conclude with a preliminary discussion of our experience from a human bidding experiment, involving Vos planners competing for orders both against each other and against some (simple) automated strategies.

85A
Pan-supplier Stock Control in a Virtual Warehouse
Emad El-Deen El-Akehal, Julian Padget
We describe the commercial application of agents to the handling of catalogue and stock-control for the selling of books on the internet. The primary characteristic of the target market is (very) low volumes over a (very) large number of items, thus agility and extremely low overheads are the essential factors for a viable business model. Being a new company (established 2004), without legacy software and with the freedom to make new choices, it was decided that the agent abstraction offered both short-term software engineering and longer-term business advantages. This expectation has been borne out in practice, in that it has been possible to construct an e-trading platform, using a 4-person team over a period of a few months, and that is now part of a live business operation handling just over 12,000 transactions daily. In this paper we explain how agents helped focus attention on the responsibilities of key software functions, how different functions should interact with one another and how to identify and propagate key performance indicator information through the system to detect unexpected behaviour. Agent technology has many potential benefits for dynamic fast-moving businesses where software requirements change quickly and business needs grow rapidly, all within a dynamic environment that has entirely different rules across the axes of geography, market, customer and competitor. Using autonomous agents allowed The Book Depository to build quickly a complex network of P2P relationships with a large number of suppliers and publishers of very different sizes who each utilize a variety of different trading and data interchange standards.

****88A****
WADE: A software platform to develop mission critical applications exploiting agents and workflows
Giovanni Caire, Danilo Gotta, Massimo Banzi
In this paper, we describe two mission critical applications currently deployed by Telecom Italia in the Operations Support System domains. The first one called "Network Neutral Element Manager" implements a mediation layer between network elements and OSS systems. The second one, known as "Wizard", provides step-by-step guidance to technicians performing maintenance operations in the fields. Both applications have strong requirements in terms of scalability and flexibility and exploit the combination of agents and workflows to meet them. As such both of them are based on a common software platform called WADE (Workflows and Agents Development Environment). WADE is the main evolution of JADE a popular Open Source framework that facilitates the development of interoperable intelligent multi-agent systems. WADE adds to JADE the support for the execution of tasks defined according to the workflow metaphor and a number of mechanisms that help managing the complexity of the distribution both in terms of administration and fault tolerance. In this paper in particular we focus on the workflow aspect and we show how WADE tries to bring the workflow approach from the business process level to the level of system internal logic.

80B
BDI-Agents for Agile Goal-Oriented Business Processes
Birgit Burmeister, M. Arnold, Felicia Copaciu, Giovanni Rimassa
Business processes are the core assets of enterprises. They turn the business potential into actual competitiveness on the market. To face the challenges posed by today's changing and uncertain business environment, traditional business process management (BPM) approaches are not sufficient anymore. This paper presents an approach to business process management, which leverages Agent Technology, especially BDI-Agent features to obtain agile business process behavior. This paper sketches the problem, describes the solution approach, and presents the experiences gained in a concrete case study in the domain of Engineering Change Management.

83B
Agent-based Patient Admission Scheduling in Hospitals
Anke K. Hutzschenreuter, Peter A. N. Bosman, Ilona Blonk-Altena, Jan van Aarle, Han La Poutré
Scheduling decisions in hospitals are often taken in a decentralized way. This means that different specialized hospital units decide autonomously on patient admissions

or operating room schedules. In this paper we present an agent-based model for the selection of an optimal mix for patient admissions. Admitting the right mix of patients is important in order to optimize the resource usage and patient throughput. Our model is based on an extensive case analysis, involving data analysis and interviews, conducted in a case study at a large hospital in the Netherlands. We focus on the coordination of different surgical patient types with probabilistic treatment processes involving multiple hospital units. We also consider the unplanned arrival of other patients requiring (partly) the same hospital resources. Simulation experiments show the applicability of our agent-based decision support tool. The simulation tool allows for the assessment of resource network usage as a function of different policies for decision making. Furthermore, the tool incorporates a first optimization module for the resource allocation of postoperative care beds.

Economic Paradigms IV

Chair: Vincent Conitzer

223A

Automated Design of Scoring Rules by Learning from Examples

Ariel Procaccia, Aviv Zohar, Jeffrey Rosenschein

Scoring rules are a broad and concisely-representable class of voting rules which includes, for example, Plurality and Borda. Our main result asserts that the class of scoring rules, as functions from preferences into candidates, is efficiently learnable in the PAC model. We discuss the applications of this result to automated design of scoring rules. We also investigate possible extensions of our approach, and (along the way) we establish a lemma of independent interest regarding the number of distinct scoring rules.

231C

Evaluation of Election Outcomes under Uncertainty

Noam Hazon, Yonatan Aumann, Sarit Kraus, Michael Wooldridge

We investigate the extent to which it is possible to evaluate the probability of a particular candidate winning an election, given imperfect information about the preferences of the electorate. We assume that for each voter, we have a probability distribution over a set of preference orderings. Thus, for each voter, we have a number of possible preference orderings – we do not know which of these orderings actually represents the voters' preferences, but we know for each one the probability that it does. We give a polynomial algorithm to solve the problem of computing the probability that a given candidate will win when the number of candidates is a constant. However, when the number of candidates is not bounded, we prove that the problem becomes #P-Hard for the Plurality, Borda, and Copeland voting protocols. We further show that even evaluating if a candidate has any chance to win is NP-Complete for the Plurality voting protocol, in the weighted voters case. We give a polynomial algorithm for this problem when the voters' weights are equal.

239B

Complexity of Terminating Preference Elicitation

Toby Walsh

Elections are a natural mechanism to aggregate agents' preferences. A number of results have shown that complexity theory is a useful tool to study computational issues surrounding the elicitation of agents' preferences, as well as the strategic manipulation of elections aggregating together agents' preferences. We study here the complexity of determining when we can terminate eliciting preferences, and prove that the complexity depends on the elicitation strategy. We show, for instance, that it may be better from a computational perspective to elicit all preferences from one agent at a time than to elicit individual preferences from multiple agents. We also study the connection between the strategic manipulation of an election and preference elicitation. We show that what we can manipulate affects the computational complexity of manipulation. In particular, we prove that there are voting rules which are easy to manipulate if we can change all of an agent's vote, but computationally intractable if we can change only some of their preferences. This suggests that, as with preference elicitation, a fine-grained view of manipulation may be informative. Finally, we study the connection between predicting the winner of an election and preference elicitation. Based on this connection, we identify a new voting rule where it is computationally difficult to decide the probability of a candidate winning given a probability distribution over the votes.

245B

Divide and Conquer: False-Name Manipulations in Weighted Voting Games

Yoram Bachrach, Edith Elkind

In this paper, we study false-name manipulations in weighted voting games. Weighted voting is a well-known model of cooperation among agents in decision-making domains. In such games, each of the players has a weight, and a coalition of players wins the game if its total weight exceeds a certain quota. While a player's ability to influence the outcome of the game is related to its weight, it is not always directly proportional to it. This observation has led to the concept of a power index, which is a measure of an agent's ``real power'' in this domain. One prominent power index is the Shapley-Shubik index, which has been widely used to analyze political power. This index is equal to the Shapley value of the player in the game. If an agent can alter the game so that his Shapley-Shubik index increases, this will mean that he has gained power in the game. Moreover, the Shapley value is often used to distribute the gains of the grand coalition. In this case, this alteration will also increase the agent's payoffs.
One way in which an agent can change the game (and hence his payoffs) is by distributing his weight among several false identities. We call this behavior a false-name manipulation. We show that such manipulations can indeed increase an agent's power, as determined by the Shapley-Shubik power index, or his payoffs, as given by the Shapley value. We provide upper and lower bounds on the effects of such manipulations. We then study this issue from the computational perspective, and show that checking whether a beneficial split exists is NP-hard. We also discuss efficient algorithms for restricted cases of this problem, as well as randomized algorithms for the general case.

249C

Copeland Voting: Ties Matter

Piotr Faliszewski, Edith Hemaspaandra, Henning Schnoor

We study the complexity of manipulation for a family of election systems derived from Copeland voting via introducing a parameter alpha that describes how ties in

head-to-head contests are valued. We show that the problem of manipulation for unweighted Copeland$^{\alpha}$ elections is NP-complete even if the size of the manipulating coalition is limited to two. Our result holds for all rational values of α such that $0 < \alpha < 1$ except for $\alpha = 0.5$. Since it is well known that manipulation via a single voter is easy for Copeland, our result is the first one where an election system originally known to be vulnerable to manipulation via a single voter is shown to be resistant to manipulation via a coalition of a constant number of voters. We also study the complexity of manipulation for Copeland$^{\alpha}$ for the case of a constant number of candidates. We show that here the exact complexity of manipulation often depends closely on the winner model as well as on the parameter alpha: Depending whether we try to make our favorite candidate a winner or a unique winner and whether alpha is 0, 1 or between these values, the problem of weighted manipulation for Copeland$^{\alpha}$ with three candidates is either in P or is NP-complete. Our results show that ways in which ties are treated in an election system, here Copeland voting, can be crucial to establishing complexity results for this system.

250A
A Broader Picture of the Complexity of Strategic Behavior in Multi-Winner Elections
Reshef Meir, Ariel Procaccia, Jeffrey Rosenschein
Recent work by Procaccia, Rosenschein and Zohar [14] established some results regarding the complexity of manipulation and control in elections with multiple winners, such as elections of an assembly or committee; this work provided an initial understanding of the topic. In this paper, we paint a much more complete picture. Indeed, we investigate four prominent multi-winner voting rules. Firstly, we characterize the complexity of manipulation and control in these voting rules under various kinds of formalizations of the manipulator's goal. Secondly, we extend the results about the complexity of control to different well-known types of control. This work enhances our comprehension of which multiwinner voting rules should be employed in various settings.

Thursday, May 15th, 13:40-15:00 – Session 5

Virtual Agents II

Chair: Elisabeth Sklar

2B
Being a Part of the Crowd: Toward Validating VR Crowds Using Presence
Nuria Pelechano, Catherine Stocker, Jan Allbeck, Norman Badler
Crowd simulation models are currently lacking a commonly accepted validation method. In this paper, we propose level of presence achieved by a human in a virtual environment (VE) as a metric for virtual crowd behavior. Using experimental evidence from the presence literature and the results of a pilot experiment that we ran, we explore the egocentric features that a crowd simulation model should have in order to achieve high levels of presence and thus be used as a framework for validation of simulated crowd behavior. We implemented four crowd models for our pilot experiment: social forces, rule based, cellular automata and HiDAC. Participants interacted with the crowd members of each model in an immersive virtual environment for the purpose of studying presence in virtual crowds, with the goal of establishing the basis for a future validation method.

5B
Simulation of Individual Spontaneous Reactive Behavior
Alejandra Garcia-Rojas, Mario Gutierrez, Daniel Thalmann
The context of this work is the search for realism and believability of Virtual Humans. Our contribution to achieve this goal is to enable Virtual Humans (VH) to react to spontaneous events in virtual environments (VE). In order to reflect the individuality of each VH, these reactions have to be expressive and unique. In this paper we present firstly a model of reaction based on personality traits. The model was defined using statistical analysis of real people reacting to unexpected events. We also consider that the emotional state is involved in the modulation of reactions, thus we integrate a model of emotion update. Secondly, we present a semantic-based methodology to compose reactive animation sequences using inverse kinematics (IK) and key frame (KF) interpolation animation techniques. Finally, we present an application that demonstrates how Virtual Humans can produce different movements as reaction to unexpected stimuli, depending on their personality traits and emotional state.

****8B****
SmartBody: Behavior Realization for Embodied Conversational Agents
Marcus Thiebaux, Andrew Marshall, Stacy Marsella, Marcelo Kallmann
Researchers demand much from their embodied conversational agents (ECAs), requiring them to be both life-like, as well as responsive to events in an interactive setting. We find that a flexible combination of animation approaches may be needed to satisfy these needs. In this paper we present SmartBody, an open source modular framework for animating ECAs in real time, based on the notion of hierarchically connected animation controllers. Controllers in SmartBody can employ arbitrary animation algorithms such as keyframe interpolation, motion capture or procedural animation. Controllers can also schedule or combine other controllers. We discuss our architecture in detail, including how we incorporate traditional approaches, and develop the notion of a controller as a reactive module within a generic framework, for realizing modular animation control. To illustrate the versatility of the architecture, we also discuss a range of applications that have used SmartBody successfully.

11B
Distinctiveness in multimodal behaviors
Maurizio Mancini, Catherine Pelachaud
While talking, people may move heavily their arms around, remain expressionless, or even display subtle facial movements... These differences may arise from personality, cultural, social factors and many more. In the present work, we are interested in defining a schema that characterizes distinctiveness in behaviors. Distinctiveness encompasses behaviors differences regarding (i) shape (which signals are performed) and (ii) quality (expressivity of movement, the way in which movements are performed). Thus, we aim at defining embodied conversational agents (ECAs) that, given their communicative intention and behaviors tendencies definition, present distinctive behaviors. Our algorithm

embeds two main steps: selection of the behaviors to convey a given communicative intention and production of qualitative behaviors. In this paper we particularly focus on the first part of the multimodal selection of signals.

Agent and Multi-Agent Learning II

Chair: Kagan Tumer

206B
Modeling how Humans Reason about Others with Partial Information
Sevan Ficici, Avi Pfeffer
Computer agents participate in many collaborative and competitive multiagent domains in which humans make decisions. For computer agents to interact successfully with people in such environments, an understanding of human reasoning is beneficial. In this paper, we investigate the question of how people reason strategically about others under uncertainty and the implications of this question for the design of computer agents. Using a situated partial-information negotiation game, we conduct human-subjects trials to obtain data on human play. We then construct a hierarchy of models that explores questions about human reasoning: Do people explicitly reason about other players in the game? If so, do people also consider the possible states of other players for which only partial information is known? Is it worth trying to capture such reasoning with computer models and subsequently utilize them in computer agents? We further address these questions by constructing computer agents that use our models. These agents utilize our models in one of two ways: one type of agent uses a model to emulate human behavior; the other type of agent uses a model of human behavior to formulate a best response. After building our agents, we deploy them in further human-subjects trials for evaluation. Our results indicate that people do reason about other players in our game and also reason under uncertainty. Better models are shown to yield more successful computer agents.

195C
Simultaneously Modeling Humans' Preferences and their Beliefs about Others' Preferences
Sevan Ficici, Avi Pfeffer
In strategic multiagent decision making, it is often the case that a strategic reasoner must hold beliefs about other agents and use these beliefs to inform its decision making. The behavior thus produced by the reasoner involves an interaction between the reasoner's beliefs about other agents and the reasoner's own preferences. A significant challenge faced by model designers, therefore, is how to model such a reasoner's behavior so that the reasoner's preferences and beliefs can each be identified and distinguished from each other. In this paper, we introduce a model of strategic reasoning that allows us to distinguish between the reasoner's utility function and the reasoner's beliefs about another agent's utility function as well as the reasoner's beliefs about how that agent might interact with yet other agents. We show that our model is uniquely identifiable. That is, no two different parameter settings will cause the model to give the same behavior over all possible inputs. We then illustrate the performance of our model in a multiagent negotiation game played by human subjects. We find that our subjects have slightly incorrect beliefs about other agents in the game.

198C
Opponent Modelling in Automated Multi-Issue Negotiation Using Bayesian Learning
Dmytro Tykhonov, Koen Hindriks
The efficiency of automated multi-issue negotiation depends on the availability and quality of knowledge about an opponent. We present a generic framework based on Bayesian learning to learn an opponent model, i.e. the issue preferences as well as the issue priorities of an opponent. The algorithm proposed is able to effectively learn opponent preferences from bid exchanges by making some assumptions about the preference structure and rationality of the bidding process. The assumptions used are general and consist among others of assumptions about the independency of issue preferences and the topology of functions that are used to model such preferences. Additionally, a rationality assumption is introduced that assumes that agents use a concession-based strategy. It thus extends and generalizes previous work on learning in negotiation by introducing a technique to learn an opponent model for multi-issue negotiations. We present experimental results demonstrating the effectiveness of our approach and discuss an approximation algorithm to ensure scalability of the learning algorithm.

201C
A Few Good Agents: Multi-Agent Social Learning
Jean Oh, Stephen Smith
In this paper, we investigate multi-agent learning (MAL) in a multi-agent resource selection problem (MARS) in which a large group of agents are competing for common resources. Since agents in such a setting are self-interested, MAL in MARS domains typically focuses on the convergence to a set of non-cooperative equilibria. As seen in the example of prisoner's dilemma, however, selfish equilibria are not necessarily optimal with respect to the natural objective function of a target problem, e.g., resource utilization in the case of MARS. Conversely, a centrally administered optimization of physically distributed agents is infeasible in many real-life applications such as transportation traffic problems. In order to explore the possibility for a middle ground solution, we analyze two types of costs for evaluating MAL algorithms in this context. The quality loss of a selfish algorithm can be quantitatively measured by the price of anarchy, i.e., the ratio of the objective function value of a selfish solution to that of an optimal solution. Analogously, we introduce the price of monarchy of a learning algorithm to quantify the practical cost of coordination in terms of communication cost. We then introduce a multi-agent social learning approach named A Few Good Agents (AFGA) that motivates self-interested agents to cooperate with one another to reduce the price of anarchy, while bounding the price of monarchy at the same time. A preliminary set of experiments on the El Farol bar problem, a simple example of MARS, show promising results.

Agent Reasoning III

Chair: Ed Durfee

145A

An Exact Algorithm for Solving MDPs under Risk-Sensitive Objective with One-Switch Utility Functions

Yaxin Liu, Sven Koenig

One-switch utility functions are an important class of nonlinear utility functions that can model humans whose decisions change with their wealth level. We study how to maximize the expected utility of Markov decision problems with given one-switch utility functions. We first utilize the fact that one-switch utility functions are weighted sums of linear and exponential utility functions to prove that there exists an optimal policy that is both stationary and deterministic as the wealth level approaches negative infinity. We then develop a solution method, the backward-induction method, that starts with this policy and then augments it for higher and higher wealth levels. Our backward-induction method determines maximal expected utilities in finite time (similar to policy iteration for linear utility functions), different from the previous functional value iteration method that typically determines only approximately maximal expected utilities (similar to value iteration for linear utility functions).

148A

Controlling Deliberation in a Markov Decision Process-Based Agent

George Alexander, Anita Raja, David Musliner

Meta-level control manages the allocation of limited resources to deliberative actions. This paper discusses efforts in adding meta-level control capabilities to a Markov Decision Process (MDP)-based scheduling agent. The agent's reasoning process involves continuous partial unrolling of the MDP state space and periodic reprioritization of the states to be expanded. The meta-level controller makes situation-specific decisions on when the agent should stop unrolling in order to derive a partial policy while bounding the costs of state reprioritization. The described approach uses performance profiling combined with multi-level strategies in its decision making. We present results showing the performance advantage of dynamic meta-level control for this complex agent.

151A

Generalized Adaptive A*

Xiaoxun Sun, Sven Koenig, William Yeoh

Agents often have to solve series of similar search problems. Adaptive A* is a recent incremental heuristic search algorithm that solves series of similar search problems faster than individual A* searches because it updates the heuristics using information from previous searches. It basically transforms consistent heuristics into more informed consistent heuristics. This allows it to find shortest paths in search spaces where the costs of actions can increase over time since consistent heuristics remain consistent after action cost increases. However, it is not guaranteed to find shortest paths in search spaces where the costs of actions can decrease over time because consistent heuristics do not necessarily remain consistent after action cost decreases. Thus, the heuristics need to get corrected explicitly after action cost decreases. In this paper, we show how to do that, resulting in a generalized version of Adaptive A* that finds shortest paths in search spaces where the costs of actions can increase or decrease over time. Our experiments demonstrate that Generalized Adaptive A* outperforms both individual A* searches and D* Lite for moving-target search, where D* Lite is an alternative state-of-the-art incremental heuristic search algorithm that finds shortest paths in search spaces where the costs of actions can increase or decrease over time.

154A

A Model of Contingent Planning for Agent Programming Languages

Yves Lesperance, Giuseppe De Giacomo, Atalay Ozgovde

In this paper, we develop a formal model of planning for an agent that is operating in a dynamic and incompletely known environment. We assume that both the agent's task and the behavior of the agents in the environment are expressed as high-level nondeterministic concurrent programs in some agent programming language (APL). In this context, planning must produce a deterministic conditional plan for the agent that can be successfully executed against all possible executions of the environment program. We handle actions with nondeterministic effects, as well as sensing actions, by treating them as actions that trigger an environmental reaction that is not under the planning agent's control. Our model of contingent planning is specified for a generic APL with a transition semantics. Within this model, we devise a general procedure for computing the contingent plans. We also show how the model can instantiated in the Situation Calculus with programs for the agent and the environment expressed in ConGolog, and we describe an implementation of the planning mechanism in this case.

Industry and Applications Track II

Chair: Michael Berger

81C

Case Studies for Contract-based Systems

Michal Jakob, Michal Pěchouček, Simon Miles, Michael Luck

Of the ways in which agent behaviour can be regulated in a multiagent system, electronic contracting – based on explicit representation of different parties' responsibilities, and the agreement of all parties to them – has significant potential for modern industrial applications. Based on this assumption, the CONTRACT project aims to develop and apply electronic contracting and contract-based monitoring and verification techniques in real world applications. This paper presents results from the initial phase of the project, which focused on requirements solicitation and analysis. Specifically, we survey four use cases from diverse industrial applications, examine how they can benefit from an agent-based electronic contracting infrastructure and outline the technical requirements that would be placed on such an infrastructure. We present the designed CONTRACT architecture and describe how it may fulfil these requirements. In addition to motivating our work on the contractbased infrastructure, the paper aims to provide a much needed community resource in terms of use case themselves and to provide a clear commercial context for the development of work on contract-based system.

84C

Electronic contracting in aircraft aftercare: A case study

Felipe Meneguzzi, Simon Miles, Michael Luck, Camden Holt, Malcolm Smith

Distributed systems comprised of autonomous self-interested entities require some sort of control mechanism to ensure the predictability of the interactions that drive

them. This is certainly true in the aerospace domain, where manufacturers, suppliers and operators must coordinate their activities to maximise safety and profit, for example. To address this need, the notion of norms has been proposed which, when incorporated into formal electronic documents, allow for the specification and deployment of contractdriven systems. In this context, we describe the CONTRACT framework and architecture for exactly this purpose, and describe a concrete instantiation of this architecture as a prototype system applied to an aerospace aftercare scenario.

86B

A Multi-Agent Simulation System for Prediction and Scheduling of Aero Engine Overhaul

Armin Stranjak, Partha Sarathi Dutta, Mark Ebden, Alex Rogers, Perukrishnen Vytelingum

The Aero Repair and Overhaul industry is facing an increasing challenge of prediction and scheduling of engine overhauls to remain competitive in a complex business arena. An appropriate technology solution is required to achieve effcient schedules while satisfying multiple opposing constraints in a highly dynamic environment. In this paper, we describe *Overhaul Prediction and Scheduling*, an agent-based simulator developed to tackle this challenge. Using negotiation strategies, it deals with the multi-dimensional scheduling optimisation problem by trading off repair costs, capacity and capability of overhaul bases, among others, in light of in-service unforseen events. It supports effective strategic decision-making via business scenario modelling.

89B

Transitioning Multiagent Technology to UAV Applications

Paul Scerri, Tracy Von Gonten, Gerald Fudge, Sean Owens, Katia Sycara

This paper describes the transition of academically developed multiagent technology for UAV coordination to an industrially developed application. The specific application is the use of lightweight UAVs with small Received Signal Strength Indicator sensors to cooperatively locate targets emitting radio frequency signals in a large area. It is shown that general techniques can be effectively transitioned, sometimes with minimal changes. However, clear differences in engineering and testing requirements of academia and commercialization require extensive effort in developing simulation and live flight testbeds. Although the technology has not yet been commercialized, initial live flight testing shows the potential of the approach.

Economic Paradigms V

Chair: Makoto Yokoo

224B

Power and Stability in Connectivity Games

Yoram Bachrach, Jeffrey Rosenschein, Ely Porat

We consider computational aspects of a game theoretic approach to network reliability. Consider a network where failure of one node may disrupt communication between two other nodes. We model this network as a simple coalitional game, called the vertex Connectivity Game (CG). In this game, each agent owns a vertex, and controls all the edges going to and from that vertex. A coalition of agents wins if it fully connects a certain subset of vertices in the graph, called the *primary* vertices.

We show that power indices, which express an agent's ability to affect the outcome of the vertex connectivity game, can be used to identify significant possible points of failure in the communication network, and can thus be used to increase network reliability. We show that in general graphs, calculating the Banzhaf power index is #P-complete, but suggest a polynomial algorithm for calculating this index in trees. We also show a polynomial algorithm for computing the core of a CG, which allows a stable division of payments to coalition agents.

232A

A Tractable and Expressive Class of Marginal Contribution Nets and Its Applications

Edith Elkind, Leslie Ann Goldberg, Paul Goldberg, Michael Wooldridge

Coalitional games raise a number of important questions from the point of view of computer science, key among them being how to represent them compactly, and how to efficiently compute solution concepts assuming such representations. *Marginal contribution nets* (MC-nets), introduced by Ieong and Shoham, are one of the simplest and most influential representation schemes for coalitional games. MC-nets are a rule-based formalism, in which rules take the form *pattern* $\rightarrow$ *value*, where "*pattern*" is a Boolean condition over agents, and "*value*" is a numeric value. Ieong and Shoham showed that, for a class of what we will call "basic" MC-nets, where patterns are constrained to be a conjunction of literals, marginal contribution nets permit the easy computation of solution concepts such as the Shapley value. However, there are very natural classes of coalitional game that require an exponential number of such basic MC-net rules. We present *read-once MC-nets*, a new class of MC-nets that is provably more compact than basic MC-nets, while retaining the attractive computational properties of basic MC-nets. We show how the techniques we develop for read-once MC-nets can be applied to other domains, in particular, computing solution concepts in network flow games on series-parallel networks.

****240C****

Cooperative Boolean Games

Paul Dunne, Sarit Kraus, Wiebe van der Hoek, Michael Wooldridge

We present and formally investigate *Cooperative Boolean Games*, a new, natural family of coalitional games that are both compact and expressive. In such a game, each agent desires to accomplish its personal goal, represented as a propositional (Boolean) formula over some set of Boolean variables; each agent is assumed to exercise unique control over some subset of the overall set of Boolean variables; the set of valuations to these variables corresponds to the set of actions the agent can take. However, the actions available to an agent are assumed to have some cost, and an agent's secondary goal is to minimise its costs. Typically, an agent must cooperate with others because it does not have sufficient control to ensure its goal is satisfied. However, the desire to minimise costs leads to preferences over possible coalitions, and hence to strategic behaviour. Following an introduction to the formal framework of Cooperative Boolean Games, we investigate three solution concepts for them: the core, stable sets, and bargaining solutions.

In each case, we characterise the complexity of the associated solution concept, and discuss the surrounding issues.

246C
Coalitional Skill Games
Yoram Bachrach, Jeffrey Rosenschein
We consider *Coalitional Skill Games (CSGs)*, a simple model of cooperation among agents. This is a restricted form of coalitional games, where each agent has a set of skills that are required to complete various tasks. Each task requires a set of skills in order to be completed, and a coalition can accomplish the task only if the coalition's agents cover the set of required skills for the task. The gain for a coalition depends only on the subset of tasks it can complete.
We consider the computational complexity of several problems in CSGs, for example, testing if an agent is dummy or a veto agent, computing the core of the game or testing whether the core is empty, and finding the Shapley value or Banzhaf power index of agents.

Thursday, May 15th, 16:20-17:20 – Session 6

Multi-Robotics II

Chair: Paul Scerri

37A
Autonomous Geocaching: Navigation and Goal Finding in Outdoor Domains
James Neufeld, Michael Bowling, Jason Roberts, Stephen Walsh, Adam Milstein, Michael Sokolsky
This paper describes an autonomous robot system designed to solve the challenging task of geocaching. Geocaching involves locating a goal object in an outdoor environment given only its rough GPS position. No additional information about the environment such as road maps, waypoints, or obstacle descriptions is provided, nor is there often a simple straight line path to the object. This is in contrast to much of the research in robot navigation which often focuses on common structural features, e.g., road following, curb avoidance, or indoor navigation. In addition, uncertainty in GPS positions requires a final local search of the target area after completing the challenging navigation problem. We describe a relatively simple robotic system for completing this task. This system addresses three main issues: building a map from raw sensor readings, navigating to the target region, and searching for the target object. We demonstrate the effectiveness of this system in a variety of complex outdoor environments and compare our system's performance to that of a human expert teleoperating the robot.

40A
The Impact of Adversarial Knowledge on Adversarial Planning in Perimeter Patrol
Noa Agmon, Vladimir Sadov, Sarit Kraus, Gal Kaminka
This paper considers the problem of multi-robot patrolling around a closed area, in the presence of an adversary trying to penetrate the area. Previous work on planning in similar adversarial environments addressed worst-case settings, in which the adversary has full knowledge of the defending robots. It was shown that non deterministic algorithms may be effectively used to maximize the chances of blocking such a full-knowledge opponent, and such algorithms guarantee a "lower bound" to the performance of the team. However, an open question remains as to the impact of the knowledge of the opponent on the performance of the robots. This paper explores this question in depth and provides theoretical results, supported by extensive experiments with *68* human subjects concerning the compatibility of algorithms to the extent of information possessed by the subjects. First, we analytically examine the case of a zero-knowledge opponent – a different extreme – and show that surprisingly, this seemingly best-case scenario (from the point of view of defending robots) is optimally addressed by a deterministic, non-randomizing patrol. Moreover, we show empirically that an optimal algorithm for the full-knowledge opponent fails miserably in this case. We then address the case in which the adversary gained partial information, propose the combine algorithm that maximizes the expected probability of penetration detection along with minimizing the deviation between the probabilities of penetration detection along the perimeter, and support the performance of this algorithm in the experiments.

43A
A Realistic Model of Frequency-Based Multi-Robot Fence Patrolling
Yehuda Elmaliach, Asaf Shiloni, Gal Kaminka
There is growing interest in multi-robot frequency-based patrolling, in which a team of robots optimizes its frequency of point visits, for every point in a target work area. In particular, recent work on patrolling of open polygons (e.g., open-ended fences) has proposed a general cooperative patrolling algorithm, in which robots move back and forth along the polygon, in an synchronized manner, such that their assigned areas of movement overlap. If the overlap factor is carefully chosen---based on the motion models of the robots---specific performance criteria are optimized. Unfortunately, previous work has presented analysis of motion models in which there are no errors in the movement of the robots, and no velocity changes. We go a step beyond existing work, and develop a realistic model of robot motion, that considers velocity uncertainties. We mathematically analyze the model and show how to use it to find optimal patrolling parameters, given known bounds of uncertainty on the motion. We then use the model to analyze the independently-programmed patrolling movements of physical robots, in extensive experiments. We show that the model predicts the behavior of the robots much more accurately than previously-described models.

Agent-based System Development III

Chair: Takayuki Ito

57C
A Model-driven, Agent-based Approach for the Integration of Services into a Collaborative BP
Ingo Zinnikus, Christian Hahn, Klaus Fischer
In cross-organisational business interactions, the most desirable solution for integrating different partners would suggest to integrate their processes and data on a rather low level. However, the internal processes and interfaces of the participating partners are often pre-existing and have to be taken as given. Furthermore, in cross-organisational scenarios partners are typically very sensitive about their product data and the algorithms that process it. In many cases, private processes are only partially visible and hidden behind public interface

descriptions. This imposes restrictions on the possible solutions for the problems which occur when partner processes are integrated. In this paper, we describe a solution which supports rapid prototyping by combining a model-driven framework for cross-organisational business processes with an agent-based approach for flexible process execution. We show how the model-driven approach can be combined with semantic Web service discovery for flexible service composition.

60C
Flexible Service Provisioning with Advance Agreements
Sebastian Stein, Terry Payne, Nick Jennings
In this paper, we propose a novel algorithm that allows service consumers to automatically select and provision service providers for their workflows in highly dynamic and uncertain multi-agent systems, where service provision is negotiated in advance using a market-based mechanism. In contrast to existing work, our algorithm reasons explicitly about the impact of failures on the overall feasibility of a workflow, and it mitigates them by proactively provisioning multiple providers in parallel for particularly critical tasks and by explicitly planning for contingencies. Furthermore, our algorithm provisions only part of its workflow at any given time, in order to retain flexibility and to decrease the potential for missing negotiated service time slots. We show empirically that current approaches are unable to achieve a high utility in uncertain and dynamic environments; whereas our proposed algorithm consistently outperforms them over a range of environments. Specifically, our approach achieves an up to 27-fold increase in utility and successfully completes most workflows within a strict deadline, even when the majority of providers do not honour their service contracts.

63C
Exploiting Organisational Information for Service Coordination in Multiagent Systems
Alberto Fernandez, Sascha Ossowski
Service-Oriented Computing and agent technology are nowadays two of the most active research areas in distributed and open systems. However, when trying to bridge the two worlds, it becomes apparent that the interaction-centric approach of multiagent systems may affect the way services are modelled and enacted, and vice versa. We claim that organisational models that underlie multiagent interactions are crucial in order to take advantage of this interrelation. In this paper we present an approach for modelling organisational structures in service-oriented multiagent systems, and show how it affects semantic service descriptions. We also present approaches to service matchmaking as well as to service composition, capable of exploiting organisational information in service descriptions. In both cases, we prove experimentally the validity of our approach.

Agent Cooperation III

Chair: Chengqi Zhang

169A
BnB-ADOPT: An Asynchronous Branch-and-Bound DCOP Algorithm
William Yeoh, Ariel Felner, Sven Koenig
Distributed constraint optimization problems (DCOPs) are a popular way of formulating and solving agent-coordination problems. It is often desirable to solve DCOPs optimally with memory-bounded and asynchronous algorithms. We thus introduce Branch-and-Bound ADOPT (BnB-ADOPT), a memory-bounded asynchronous DCOP algorithm that uses the message passing and communication framework of ADOPT, a well known memory-bounded asynchronous DCOP algorithm, but changes the search strategy of ADOPT from best-first search to depth-first branch-and-bound search. Our experimental results show that BnB-ADOPT is up to one order of magnitude faster than ADOPT on a variety of large DCOPs and faster than NCBB, a memory-bounded synchronous DCOP algorithm, on most of these DCOPs.

172A
Evaluating the Performance of DCOP Algorithms in a Real World, Dynamic Problem
Robert Junges, Ana Bazzan
Complete algorithms have been proposed to solve problems modelled as distributed constraint optimization (DCOP). However, there are only few attempts to address real world scenarios using this formalism, mainly because of the complexity associated with those algorithms. In the present work we compare three complete algorithms for DCOP, aiming at studying how they perform in complex and dynamic scenarios of increasing sizes. In order to assess their performance we measure not only standard quantities such as number of cycles to arrive to a solution, size and quantity of exchanged messages, but also computing time and quality of the solution which is related to the particular domain we use. This study can shed light in the issues of how the algorithms perform when applied to problems other than those reported in the literature (graph coloring, meeting scheduling, and distributed sensor network).

175A
On K-Optimal Distributed Constraint Optimization Algorithms: New Bounds and Algorithms
Emma Bowring, Jonathan Pearce, Christopher Portway, Manish Jain, Milind Tambe
Distributed constraint optimization (DCOP) is a promising approach to coordination, scheduling and task allocation in multiagent networks. In large-scale or low-bandwidth networks, finding the global optimal is often impractical. K-optimal algorithms is a promising new approach: for the first time they provide us a set of locally optimal algorithms with quality guarantees. Unfortunately, previous work in k-optimality did not address domains where we may have prior knowledge of reward structure; and it failed to provide quality guarantees or algorithms for domains with hard constraints (such as agents' local resource constraints). This paper addresses these shortcomings with three key contributions. It provides: (i) improved lower-bounds on k-optima quality incorporating available prior knowledge of reward structure; (ii) lower bounds on k-optima quality for problems with hard constraints; and (iii) k-optimal algorithms for solving DCOPs with hard constraints and detailed experimental results on large-scale networks.

Industry and Applications Track III

Chair: Satoshi Nishiyama

87C

Autonomic Multi-Agent Management of Power and Performance in Data Centers

Rajarshi Das, Jeffrey O. Kephart, Charles Lefurgy, Gerald Tesauro, David W. Levine, Hoi Chan

The rapidly rising cost and environmental impact of energy consumption in data centers has become a multi-billion dollar concern globally. In response, the IT Industry is actively engaged in a first-to-market race to develop energy-conserving hardware and software solutions that do not sacrifice performance objectives. In this work we demonstrate a prototype of an integrated data center power management solution that employs server management tools, appropriate sensors and monitors, and an agent-based approach to achieve specified power and performance objectives. By intelligently turning off servers under low-load conditions, we can achieve over 25% power savings over the unmanaged case without incurring SLA penalties for typical daily and weekly periodic demands seen in webserver farms.

****88A****

Cooperative search for optimizing pipeline operations

T. Mora, A.B. Sesay, J. Denzinger, H. Golshan, G. Poissant, C. Konecnik

We present an application of a multi-agent cooperative search approach to the problem of optimizing gas pipeline operations, i.e. finding control parameters for a gas transmission network that result in a low usage of energy to make the required gas deliveries. Our cooperative search approach improves on the pure competition of search agents by having them exchange good solutions from time to time that both are integrated into the search state of the agents and used to improve the search control of the agents. Our experimental evaluation with real problem instances from TransCanada show that our system meets TransCanada's time requirements and reliably outperforms the interactive method that is the current state-of-the-art by creating solutions that require more than 10 percent less energy.

****79A****

Deployed ARMOR Protection: The Application of a Game Theoretic Model for Security at the Los Angeles International Airport

James Pita, Manish Jain, Janusz Marecki, Fernando Ordóñez, Christopher Portway, Milind Tambe, Craig Western, Praveen Paruchuri, Sarit Kraus

Security at major locations of economic or political importance is a key concern around the world, particularly given the threat of terrorism. Limited security resources prevent full security coverage at all times, which allows adversaries to observe and exploit patterns in selective patrolling or monitoring, e.g. they can plan an attack avoiding existing patrols. Hence, randomized patrolling or monitoring is important, but randomization must provide distinct weights to different actions based on their complex costs and benefits. To this end, this paper describes a promising transition of the latest in multi-agent algorithms – in fact, an algorithm that represents a culmination of research presented at AAMAS – into a deployed application. In particular, it describes a software assistant agent called ARMOR (Assistant for Randomized Monitoring over Routes) that casts this patrolling/monitoring problem as a Bayesian Stackelberg game, allowing the agent to appropriately weigh the different actions in randomization, as well as uncertainty over adversary types. ARMOR combines three key features: (i) It uses the fastest known solver for Bayesian Stackelberg games called DOBSS, where the dominant mixed strategies enable randomization; (ii) Its mixed-initiative based interface allows users to occasionally adjust or override the automated schedule based on their local constraints; (iii) It alerts the users if mixed-initiative overrides appear to degrade the overall desired randomization. ARMOR has been successfully deployed since August 2007 at the Los Angeles International Airport (LAX) to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals. This paper examines the information, design choices, challenges, and evaluation that went into designing ARMOR.

Economic Paradigms VI

Chair: Ulle Endriss

225C

Mechanism Design for Abstract Argumentation

Iyad Rahwan, Kate Larson

Since their introduction by Dung over a decade ago, abstract argumentation frameworks have received increasing interest in artificial intelligence as a convenient model for reasoning about general characteristics of argument. Such a framework consists of a set of arguments and a binary defeat relation among them. Various semantic and computational approaches have been developed to characterise the acceptability of individual arguments in a given argumentation framework. However, little work exists on understanding the strategic aspects of abstract argumentation among self-interested agents. In this paper, we introduce (game-theoretic) argumentation mechanism design (ArgMD), which enables the design and analysis of argumentation mechanisms for self-interested agents. We define the notion of a direct-revelation argumentation mechanism, in which agents must decide which arguments to reveal simultaneously. We then design a particular direct argumentation mechanism and prove that it is strategy proof under specific conditions; that is, the strategy profile in which each agent reveals its arguments truthfully is a dominant strategy equilibrium.

233B

Undominated VCG Redistribution Mechanisms

Mingyu Guo, Vincent Conitzer

Many important problems in multiagent systems can be seen as resource allocation problems. For such problems, the well-known Vickrey-Clarke-Groves (VCG) mechanism is efficient, incentive compatible, individually rational, and does not incur a deficit. However, the VCG mechanism is not (strongly) budget balanced: generally, the agents' payments will sum to more than 0. Very recently, several mechanisms have been proposed that redistribute a significant percentage of the VCG payments back to the agents while maintaining the other properties. This increases the agents' utilities. One redistribution mechanism dominates another if it always redistributes at least as much to each agent (and sometimes more). In this paper, we provide a characterization of undominated redistribution mechanisms. We also propose several techniques that take a dominated redistribution mechanism as input, and produce as output another redistribution mechanism that dominates the original. One technique immediately produces an undominated redistribution mechanism that is not necessarily anonymous. Another technique

preserves anonymity, and repeated application results in an undominated redistribution mechanism in the limit. We show experimentally that these techniques improve the known redistribution mechanisms.

241A

Optimal-in-Expectation Redistribution Mechanisms

Mingyu Guo, Vincent Conitzer

Many important problems in multiagent systems involve the allocation of multiple resources to multiple agents. If agents are self-interested, they will lie about their valuations for the resources if they perceive this to be in their interest. The well-known VCG mechanism allocates the items efficiently, is incentive compatible (agents have no incentive to lie), and never runs a deficit. Nevertheless, the agents may have to make large payments to a party outside the system of agents, leading to decreased utility for the agents. Recent work has investigated the possibility of redistributing some of the payments back to the agents, without violating the other desirable properties of the VCG mechanism.

We study multi-unit auctions with unit demand, for which previously a mechanism has been found that maximizes the worst-case redistribution percentage. In contrast, we assume that a prior distribution over the agents' valuations is available, and try to maximize the expected total redistribution. We analytically solve for a mechanism that is optimal among linear redistribution mechanisms. The optimal linear mechanism is asymptotically optimal. We also propose discretization redistribution mechanisms. We show how to automatically solve for the optimal discretization redistribution mechanism for a given discretization step size, and show that the resulting mechanisms converge to optimality as the step size goes to zero. We also present experimental results showing that for auctions with many bidders, the optimal linear redistribution mechanism redistributes almost everything, whereas for auctions with few bidders, we can solve for the optimal discretization redistribution mechanism with a very small step size.

Thursday, May 15th, 17:30-18:30
ACM-SIGART Award Talk

Computer Science and Game Theory

Yoav Shoham

Game theory has been playing an increasingly visible role in computer science, in areas as diverse as artificial intelligence, theory, distributed systems, and other areas. I take stock of where most of the action has been in the past decade or so, and suggest that going forward, the most dramatic interaction between computer science and game theory could be around what might be called game theory pragmatics.

Friday, May 16th, 8:40-9:40 – Invited Talk

Cooperative Control of Small and Micro Air Vehicles

Randal W. Beard

The focus of this talk will be cooperative control techniques for small and micro air vehicles. There are numerous potential applications of this technology including aerial reconnaissance, border patrol, monitoring forest fires, oil fields, and pipelines, and tracking wildlife. This talk will present a general approach to cooperative control problems which can be summarized in four steps. The first step is to identify the coordination variables which are the minimal information needed to effect cooperation. The second step is to quantify the cooperation constraint and cooperation objective in terms of the coordinate variables. The third step is to develop a centralized cooperation strategy that acts upon the instantaneous values of the coordination variables to achieve the objectives. Finally, the fourth step is to use information consensus schemes to transform the centralized strategy into a decentralized algorithm. The consensus algorithms allow a team of vehicles to agree upon the instantaneous value of the coordination variables while connected through a noisy, intermittent, time-varying communication network.

We will also show several applications of our approach in the context of small unmanned air vehicles (UAVs). The first application will be to the problem of cooperative rendezvous. There are numerous military scenarios where it is desirable to have a team of UAVs converges simultaneously to a region of interest. However, pop-up threats, wind, and an unreliable communication environment make this problem extremely challenging for small UAVs. The second application will be distributed fire perimeter monitoring where a team of UAVs is tasked to monitor the border of a forest fire. The nature of the environment only allows communication when the UAVs are in close vicinity of each other. Therefore, cooperation must be achieved with only infrequent communication. The third application will be that of cooperatively tracking and targeting a moving ground target when sensor occlusions are probable.

Friday, May 16th, 10:10-12:10 – Session 7

Virtual Agents III

Chair: Zsofia Ruttkay

****3C****

Modeling Parallel and Reactive Empathy in Virtual Agents: An Inductive Approach

Scott McQuiggan, Jennifer Robison, Robert Phillips, James Lester

Humans continuously assess one another's situational context, modify their own affective state, and then respond based on these outcomes through empathetic expression. Virtual agents should be capable of similarly empathizing with users in interactive environments. A key challenge posed by empathetic reasoning in virtual agents is determining whether to respond with parallel or reactive empathy. Parallel empathy refers to mere replication of another's affective state, whereas reactive empathy exhibits greater cognitive awareness and may lead to incongruent emotional responses (i.e., emotions different from the recipient's and perhaps intended to alter negative affect). Because empathy is not yet sufficiently well understood, it is unclear as to which type of empathy is most effective in which social situations. Devising empirically informed models of empathy from observations of "empathy in action" may lead to virtual agents that can accurately respond in social situations.

This paper proposes a unified data-driven framework for modeling parallel and reactive empathy. First, in training sessions, a trainer guides a virtual agent through a series of problem-solving tasks in a learning environment encountering empathetic characters. The proposed inductive architecture tracks situational data including actions, visited locations, intentions, and the trainer's physiological response to generate models of empathy. Empathy models are used to drive runtime situation-appropriate empathetic behaviors by selecting suitable

parallel or reactive empathetic expressions. The inductive approach has been empirically evaluated in an interactive learning environment showing that the induced empathy models are able to accurately assess social contexts and generate appropriate empathetic responses for virtual agent control.

6C
Towards Background Emotion Modeling for Embodied Virtual Agents
Luís Morgado, Graca Gaspar
For the realistic simulation of embodied agents we need a model of emotion that represents both structural and dynamic aspects of emotional phenomena to serve as background support for multifaceted emotion characterization. In this paper we present an emotion model oriented towards that aim, which provides a continuous modeling of the evolution of emotional phenomena. We also illustrate how it can be used to provide different perspectives of an emotional situation, namely by identifying emotional patterns that can be characterized as discrete emotional states.

9C
MADeM: a multi-modal decision making for social MAS
Francisco Grimaldo, Miguel Lozano Ibañez, Fernando Barber
This paper presents MADeM, a multi-modal agent decision making to provide virtual agents with socially acceptable decisions. We consider multi-modal decisions as those that are able to merge multiple information sources received from a MAS. MADeM performs social decisions as it relies on auctions, a well known market based coordination mechanism. Our social agents express their preferences, over the different solutions considered for a specific decision problem, using utility functions. Therefore, coordinated social behaviors such as task passing or planned meetings can be evaluated to finally obtain socially acceptable behaviors. Additionally, MADeM is able to simulate different kind of societies (e.g. elitist, utilitarian, etc), as well as social attitudes of their members such as, egoism, altruism, indifference or reciprocity. MADeM agents have been successfully verified in a 3D dynamic environment while simulating a virtual university bar, where different types of waiters (eg. coordinated, social, egalitarian) and customers (e.g. social, lazy) interact to finally animate complex social scenes.

12C
Dynamic Bayesian Network Based Interest Estimation for Visual Attentive Presentation Agents
Boris Brandherm, Helmut Prendinger, Mitsuru Ishizuka
In this paper, we report on an interactive system and the results of a formal user study that was carried out with the aim of comparing two approaches to estimating users' interest in a multimodal presentation based on their eye gaze. The scenario consists of a virtual showroom where two 3D agents present product items in an entertaining way, and adapt their performance according to users' (in)attentiveness. In order to infer users' attention and visual interest with regard to interface objects, our system analyzes eye movements in real-time. Interest detection algorithms used in previous research determine an object of interest based on the time that eye gaze dwells on that object. However, this kind of algorithm does not seem to be well suited for dynamic presentations where the goal is to assess the user's focus of attention with regard to a dynamically changing presentation. Here, the current context of the object of interest has to be considered, i.e., whether the visual object is part of (or contributes to) the current presentation content or not. Therefore, we recently proposed an approach that estimates the interest (or non-interest) of a user by means of dynamic Bayesian networks that take into account the current context of the attention receiving object. Hence, the presentation agents can provide timely and appropriate response. The benefits of our approach will be demonstrated both theoretically and empirically.

15C
A Model of Gaze for the Purpose of Emotional Expression in Virtual Embodied Agents
Brent Lance, Stacy Marsella
Currently, state of the art virtual agents lack the ability to display emotion as seen in actual humans, or even in hand-animated characters. One reason for the emotional inexpressiveness of virtual agents is the lack of a model of emotionally expressive gaze manner. For virtual agents to express emotion that observers can empathize with, they need the ability to generate gaze - including eye, head, and torso movement - to arbitrary targets, while displaying arbitrary emotional states. Our previous work [18] describes the Gaze Warping Transformation, a method of generating emotionally expressive head and torso movement during gaze shifts that is derived from human movement data. Through an evaluation, it was shown that applying different transformations to the same gaze shift could modify the affective state perceived when the transformed gaze shift was displayed to a human observer. In this paper we propose a model of realistic, emotionally expressive gaze that builds upon the Gaze Warping Transformation by improving the implementation of the transformations, and by adding a model of eye movement drawn from the visual neuroscience literature. We describe how to generate a gaze to an arbitrary target, while displaying an arbitrary emotional behavior. Finally, we propose an evaluation to determine what emotions human observers will attribute to the generated gaze shifts. Once this work is completed, virtual agents will have access to a new channel for emotionally expressive behavior.

18C
Politeness and Alignment in Dialogues with a Virtual Guide
Markus de Jong, Mariet Theune, Dennis Hofs
Language alignment is something that happens automatically in dialogues between human speakers. The ability to align is expected to increase the believability of virtual dialogue agents. In this paper we extend the notion of alignment to affective language use, describing a model for dynamically adapting the linguistic style of a virtual agent to the level of politeness and formality detected in the user's utterances. The model has been implemented in the Virtual Guide, an embodied conversational agent giving directions in a virtual environment. Evaluation shows that our formality model needs improvement, but that the politeness tactics used by the Guide are mostly interpreted as intended, and that the alignment to the user's language is noticeable.

Agent Reasoning IV

Chair: Helder Coelho

157A

Not All Agents Are Equal: Scaling up Distributed POMDPs for Agent Networks

Janusz Marecki, Tapana Gupta, Pradeep Varakantham, Milind Tambe, Makoto Yokoo

Many applications of networks of agents, including mobile sensor networks, unmanned air vehicles, autonomous underwater vehicles, involve 100s of agents acting collaboratively under uncertainty. Distributed Partially Observable Markov Decision Problems (Distributed POMDPs) are well-suited to address such applications, but so far, only limited scale-ups of up to five agents have been demonstrated. This paper escalates the scale-up, increasing the number of agents in distributed POMDPs for the first time into double digits. One major insight is to marry finite state machines (FSMs) for policy representation with algorithms that exploit network structures within an agent network, leading to a new algorithm called FANS. While marrying these two approaches is itself a major contribution within FANS, it also provides two additional key contributions. First, not all agents within a network need the same expressivity of policy representation; FANS introduces novel heuristics to automatically vary the FSM size in different agents for scale-up. We present detailed experimental comparison of these heuristics. Second, FANS illustrates that FSM representation enables dynamic programming for speeding up heuristic computations and policy evaluations within the network algorithms, without requiring an exponential increase in number of states as required in previous algorithms. Experimental results show not only orders of magnitude improvements over previous best known algorithms for smaller-scale domains (with similar solution quality), but also a scale-up into double digits in terms of numbers of agents.

****143B****

The Permutable POMDP: Fast Solutions to POMDPs for Preference Elicitation

Doshi Finale, Nicholas Roy

The ability for an agent to reason under uncertainty is crucial in many planning applications, since an agent rarely has access to complete, error-free information about its environment. Partially Observable Markov Decision Processes (POMDPs) are a desirable framework in these planning domains because the resulting policies allow the agent to reason about its own uncertainty. In any domain with hidden state and noisy observations, POMDPs optimally trade between actions that increase an agent's knowledge and actions that increase an agent's reward.

Unfortunately, for many real world problems, even approximating good POMDP solutions becomes computationally intractable without leveraging structure in the problem domain. We show that in particular, the structure of many preference elicitation problems---in which the agent must discover some hidden preference or desire from another (usually human) agent---allows the POMDP solution to be represented using exponentially fewer belief points than standard approaches while retaining the quality of the solution. We introduce the "Permutable POMDP'" model and show that we can generate policies for such problems more efficiently than using standard solution algorithms.

146B

Value-Based Observation Compression for DEC-POMDPs

Alan Carlin, Shlomo Zilberstein

Representing agent polices compactly is essential for improving the scalability of multi-agent planning algorithms. We focus in this paper on developing a pruning technique that allows us to eliminate certain observations from agent policies while minimizing the loss of value. This is particularly important for solving finite-horizon decentralized POMDPs, where agent polices are represented as trees. The size of policy trees grows exponentially with the number of observations. We introduce a value-based observation compression technique that prunes the least valuable observations while maintaining an error bound on the value lost as a result of pruning. We analyze the characteristics of this pruning strategy and show empirically that it is very effective. We use compact policies and obtain significantly higher values compared with the best existing DEC-POMDP algorithm.

149B

No-Regret Learning and a Mechanism for Distributed Multi-Agent Planning

Jan-P. Calliess, Geoffrey Gordon

We develop a novel mechanism for coordinated, distributed multiagent planning. We consider problems stated as a collection of single-agent planning problems coupled by common soft constraints on resource consumption. (Resources may be real or fictitious, the latter introduced as a tool for factoring the problem). A key idea is to recast the distributed planning problem as learning in a repeated game between the original agents and a newly introduced group of adversarial agents who influence prices for the resources. The adversarial agents benefit from arbitrage: that is, their incentive is to uncover violations of the resource usage constraints and, by selfishly adjusting prices, encourage the original agents to avoid plans that cause such violations. If all agents employ no-regret learning algorithms in the course of this repeated interaction, we are able to show that our mechanism can be set up to achieve design goals such as social optimality (efficiency), budget-balance and Nash-equilibrium convergence to within an error which approaches zero as the agents gain experience. In particular, the agents' average plans converge to a socially optimal solution for the original planning task. We present experiments in a simulated network routing domain demonstrating our method's ability to reliably generate sound plans.

152B

Exploiting Locality of Interaction in Factored Dec-POMDPs

Frans Oliehoek, Matthijs Spaan, Shimon Whiteson, Nikos Vlassis

Decentralized partially observable Markov decision processes (Dec-POMDPs) constitute an expressive framework for multiagent planning under uncertainty, but solving them is provably intractable. We demonstrate how their scalability can be improved by exploiting locality of interaction between agents in a factored representation. Factored Dec-POMDP representations have been proposed before, but only for Dec-POMDPs whose transition and observation models are fully independent. Such strong assumptions simplify the planning problem, but result in models with limited applicability. By contrast, we consider general factored

Dec-POMDPs for which we analyze the model dependencies over space (locality of interaction) and time (horizon of the problem). We also present a formulation of decomposable value functions. Together, our results allow us to exploit the problem structure as well as heuristics in a single framework that is based on collaborative graphical Bayesian games (CGBGs). A preliminary experiment shows a speedup of two orders of magnitude.

155B
Interaction-Driven Markov Games for Decentralized Multiagent Planning under Uncertainty
Matthijs Spaan, Francisco Melo
In this paper we propose interaction-driven Markov games (IDMGs), a new model for multiagent decision making under uncertainty. IDMGs aim at describing multiagent decision problems in which interaction among agents is a local phenomenon. To this purpose, we explicitly distinguish between situations in which agents should interact and situations in which they can afford to act independently. The agents are coupled through the joint rewards and joint transitions in the states in which they interact. The model combines several fundamental properties from transition-independent Dec-MDPs and weakly coupled MDPs while allowing to address, in several aspects, more general problems. We introduce a fast approximate solution method for planning in IDMGs, exploiting their particular structure, and we illustrate its successful application on several large multiagent tasks.

Agent Cooperation IV

Chair: Radhika Nagpal

164B
Automated Global-to-Local Programming in 1-D Spatial Multi-Agent Systems
Daniel Yamins, Radhika Nagpal
A *spatial computer* is a distributed multi-agent system that is embedded in a geometric space. A key challenge is engineering local agent interaction rules that enable spatial computers to robustly achieve global computational tasks. This paper develops a principled approach to global-to-local programming, for pattern formation problems in a one-dimensional multi-agent model. We present theoretical techniques that address the existence, construction, and resource tradeoffs of robust local rule solutions to global patterns, and combine them into a "global-to-local compiler".

167B
An Approach to Online Optimization of Heuristic Coordination Algorithms
Jumpol Polvichai, Paul Scerri, Michael Lewis
Coordination algorithms for large scale coordination are typically heuristic and hence require tuning for particular environments. In domains where characteristics of the environment can vary dramatically from scenario to scenario, it is desirable to have automated techniques for appropriately configuring automated coordination. This paper presents an approach that takes performance data from a simulator to train a stochastic neural network that concisely models the complex, probabilistic relationship between configurations, environments and performance metrics. The stochastic neural network is then used as the core of a tool that allows rapid online or offline configuration of coordination algorithms to particular scenarios and user preferences. The overall system allows automated adaptation of coordination, leading to better performance in new scenarios.

170B
Using Multi-agent Potential Fields in Real-time Strategy Games
Johan Hagelbäck, Stefan Johansson
Bots for Real Time Strategy (RTS) games provide a rich challenge to implement. A bot controls a number of units that will have to navigate in a partially unknown environment, while at the same time search for enemies and coordinate attacks to fight them down. Potential fields is a technique originating from the area of robotics where it is used in controlling the navigation of robots in dynamic environments. Although attempts have been made to transfer the technology to the gaming sector, assumed problems with efficiency and high costs for implementation have made the industry reluctant to adopt it. We present a Multi-agent Potential Field based bot architecture that is evaluated in a real time strategy game setting and compare it, both in terms of performance, and in terms of softer attributes such as configurability with other state-of-the-art solutions. Although our solution does not quite reach the performance standards of traditional RTS bots in the competition, we see great unexploited potential in using multi-agent potential field based solutions in RTS games.

173B
Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm
Alessandro Farinelli, Alex Rogers, Adrian Petcu, Nick Jennings
This paper considers the problem of performing decentralised coordination of low-power embedded devices (as is required within many environmental sensing and surveillance applications). Specifically, we address the generic problem of maximising social welfare within a group of interacting agents. We propose a novel representation of the problem as a cyclic bipartite factor graph, composed of variable and function nodes which represent the agents' states and utilities respectively. We show that such representation allows us to use an extension of the max-sum algorithm to generate approximate solutions to this global optimisation problem, through local decentralised message passing. We empirically evaluate this approach on a canonical coordination problem (graph colouring), and benchmark it against a state of the art approximate and complete algorithms (DSA and DPOP). We show that our approach is robust to lossy communication, that it generates solutions closer to those of DPOP than DSA is able to, and that it does so with a communication cost (in terms of total messages size) that scales very well with the number of agents in the system (compared to the exponential increase of DPOP). Finally, we describe a hardware implementation of our algorithm operating on low-power Chipcon CC2431 System-on-Chip sensor nodes.

176B
Look Where You Can See:Predictability & Criticality Metrics for Coordination in Complex Environments
Rajiv Maheswaran, Pedro Szekely, Marcel Becker, Stephen Fitzpatrick, Gergely Gati, Jing Jin, Robert Neches, Nader Noori, Craig Rogers, Romeo Sanchez, Kevin Smyth, Chris Van Buskirk
We address the problem of coordinating the activities of a team of agents in a dynamic, uncertain, nonlinear environment. Bounded rationality, bounded

communication, subjectivity and distribution make it extremely challenging to find effective strategies. We propose an approach, resembling the “parable of the street light”, based on Predictability and Criticality Metrics (PCM) for designing solutions in these environments. We also characterize key properties of criticality metrics that facilitate the creation of information structures and protocols needed to support this approach. In the context of the DARPA Coordinators program, we show how the PCM approach yielded a system that significantly outperformed several competing approaches in an extensive independent evaluation.

177C

Aligning social welfare and agent preferences to alleviate traffic congestion

Kagan Tumer, Zach Welch, Adrian Agogino

Multiagent coordination algorithms provide unique insights into the challenging problem of alleviating traffic congestion. What is particularly interesting in this class of problem is that no individual action (e.g., leave at a given time) is intrinsically “bad” but that combinations of actions among agents lead to undesirable outcomes. As a consequence, agents need to learn how to coordinate their actions with those of other agents, rather than learn a particular set of “good” actions. In general, the traffic problem can be approached from two distinct perspectives: (i) from a city manager's point of view, where the aim is to optimize a city wide objective function (e.g., minimize total city wide delays), and (ii) from the individual driver's point of view, where each driver is aiming to optimize a personal objective function (e.g., a ``timeliness"" function that minimizes the difference desired and actual arrival times at a destination). In many cases, these two objective functions are at odds with one another, where drivers aiming to optimize their own objectives yield to congestion and poor values of city objective functions.

In this paper we present an objective shaping approach to both types of problems and study the system behavior that arises from the drivers' choices. We first show a top-down approach that provides incentives to drivers and leads to good values of the city manager's objective function. We then present a bottom-up approach that shows that drivers aiming to optimize their own personal timeliness objective lead to poor performance with respect to a city manager's objective function. Finally, we present the intriguing result that drivers that aim to optimize a modified version of their own timeliness function not only perform well in terms of the city manager's objective function, but also perform better with respect to their own original timeliness functions.

Agent and Multi-Agent Learning III

Chair: Sandip Sen

204C

Sequential Decision Making in Repeated Coalition Formation under Uncertainty

Georgios Chalkiadakis, Craig Boutilier

The problem of coalition formation when agents are uncertain about the types or capabilities of their potential partners is a critical one. In [Chalkiadakis and Boutilier, 2004] a Bayesian reinforcement learning framework is developed for this problem when coalitions are formed (and tasks undertaken) repeatedly: not only does the model allow agents to refine their beliefs about the types of others, but uses value of information to define optimal exploration policies. However, computational approximations in that work are purely myopic. We present novel, non-myopic learning algorithms to approximate the optimal Bayesian solution, providing tractable means to ensure good sequential performance. We evaluate our algorithms in a variety of settings, and show that one, in particular, exhibits consistently good sequential performance. Further, it enables the Bayesian agents to transfer acquired knowledge among different dynamic tasks.

207C

Emerging coordination in infinite team Markov games

Francisco Melo, Isabel Ribeiro

In this paper we address the problem of coordination in infinite multiagent sequential decision problems. We adopt a game theoretic formalism to describe the interaction of the multiple decision-makers and propose the approximate biased adaptive play algorithm. This algorithm is an extension of biased adaptive play to team Markov games defined over infinite state-spaces. We establish our method to coordinate w.p.1 in the optimal strategy and discuss how this methodology can be combined with approximate learning architectures. We conclude with a simple example where the application of the algorithm is illustrated.

****196A****

Approximate Predictive State Representations

Britton Wolfe, Michael James, Satinder Singh

Predictive state representations (PSRs) are models that represent the state of a dynamical system as a set of predictions about future events. In order to learn and use PSRs to model complex dynamical systems, one must make approximations. We take the first step in developing a theory of approximate PSRs by bounding the error of the approximate predictive state under certain approximation schemes. We also introduce factored PSRs, a class of PSRs that allow one to tune the degree of approximation by trading off accuracy for compactness. We demonstrate this trade-off empirically on some example systems, using factored PSRs that were learned from data.

199A

MB-AIM-FSI: A Model Based Framework for exploiting gradient ascent MultiAgent Learners in Strategic

Doran Chakraborty, Sandip Sen

Future agent applications will increasingly represent human users autonomously or semi-autonomously in strategic interactions with similar entities. Hence, there is a growing need to develop algorithmic approaches that can learn to recognize commonalities in opponent strategies and exploit such commonalities to improve strategic response. Recently a framework has been proposed that aims for *targeted optimality* against a set of finite memory opponents. We propose an approach that aims for targeted optimality against the set of all possible multiagent learning algorithms that perform gradient search to select a single stage Nash Equilibria of a repeated game. Such opponents induce a Markov Decision Process as the learning environment and appropriate responses to such environments are learned by assuming a generative model of the environment. In the absence of a generative model, we present a framework, *MB-AIM-FSI*, that models the opponent online based on interactions, solves the model off-line

when sufficient information has been gathered, stores the strategy in the repository and finally uses it judiciously when playing against the same or similar opponent at a later time.

202A
Sigma Point Policy Iteration
Michael Bowling, Alborz Geramifard, David Wingate
In reinforcement learning, least-squares temporal difference methods (e.g., LSTD and LSPI) are effective, data-efficient techniques for policy evaluation and control with linear value function approximation. These algorithms rely on policy-dependent expectations of the transition and reward functions, which require all experience to be remembered and iterated over for each new policy evaluated. We propose to summarize experience with a compact policy-independent Gaussian model. We show how this policy-independent model can be transformed into a policy-dependent form and used to perform policy evaluation. Because closed-form transformations are rarely available, we introduce an efficient sigma point approximation. We show that the resulting Sigma-Point Policy Iteration algorithm (SPPI) is mathematically equivalent to LSPI for tabular representations and empirically demonstrate comparable performance for approximate representations. However, the experience does not need to be saved or replayed, meaning that for even moderate amounts of experience, SPPI is an order of magnitude faster than LSPI.

205A
Dynamics Based Control with PSRs
Ariel Adam, Zinovi Rabinovich, Jeffrey Rosenschein
We present an extension of the Dynamics Based Control (DBC) paradigm to environment models based on Predictive State Representations (PSRs). We show an approximate greedy version of the DBC for PSR model, *EMT-PSR*, and demonstrate how this algorithm can be applied to solve several control problems. We then provide some classifications and requirements of PSR environment models that are necessary for the EMT-PSR algorithm to operate.

Economic Paradigms VII

Chair: Jonathan Bredin

226A
Stochastic Search Methods for Nash Equilibrium Approximation in Simulation-Based Games
Yevgeniy Vorobeychik, Michael Wellman
We define the class of games called *simulation-based games*, in which the payoffs are available as an output of an oracle (simulator), rather than specified analytically or using a payoff matrix. We then describe a convergent algorithm based on a hierarchical application of simulated annealing for estimating Nash equilibria (if they exist) in simulation-based games with finite-dimensional strategy sets. Additionally, we present alternative algorithms for best response and Nash equilibrium estimation, with a particular focus on one-shot infinite games of incomplete information. Our experimental results demonstrate that all the approaches we introduce are efficacious, albeit some more so than others. We show, for example, that while iterative best response dynamics has relatively weak convergence guarantees, it outperforms our convergent method experimentally. Additionally, we provide considerable evidence that a method based on random search outperforms gradient descent in our setting.

234C
Searching for Approximate Equilibria in Empirical Games
Patrick Jordan, Yevgeniy Vorobeychik, Michael Wellman
When exploring a game over a large strategy space, it may not be feasible or cost-effective to evaluate the payoff of every relevant strategy profile. For example, determining a profile payoff for a procedurally defined game may require Monte Carlo simulation or other costly computation. Analyzing such games poses a search problem, with the goal of identifying equilibrium profiles by evaluating payoffs of candidate solutions and potential deviations from those candidates. We propose two algorithms, applicable to distinct models of the search process. In the revealed-payoff model, each search step determines the exact payoff for a designated pure-strategy profile. In the noisy-payoff model, a step draws a stochastic sample corresponding to such a payoff. We compare our algorithms to previous proposals from the literature for these two models, and demonstrate performance advantages.

242B
Computationally-efficient Winner Determination for Mixed Multi-Unit Combinatorial Auctions
Andrea Giovannucci, Meritxell Vinyals, Juan A. Rodriguez-Aguilar, Jesus Cerquides
Mixed Multi-Unit Combinatorial Auctions (MMUCAs) offer a high potential to be employed for the automated assembly of supply chains of agents offering goods and services. Its winner determination problem (WDP) is an NP-hard problem that can be mapped into an integer program. Nonetheless, the computational cost of the current solution hinders the application of MMUCAs to realistic scenarios. In this paper we propose a new integer program (IP) for MMUCAs that severely simplifies the problem by taking advantage of the topological characteristics of the WDP. Furthermore, we provide empirical evidence showing that the new IP allows to cope with much larger supply chain formation scenarios.

247B
Characterizing effective auction mechanisms: Insights from the 2007 TAC market design competition
Jinzhong Niu, Kai Cai, Enrico Gerding, Peter McBurney, Simon Parsons
This paper analyses the entrants to the 2007 TAC Market Design competition. It presents a classification of the entries to the competition, and uses this classification to compares these entries. The paper also attempts to relate market dynamics to the auction rules adopted by these entries and their adaptive strategies via a set of post-tournament experiments. Based on this analysis, the paper speculates about the design of effective auction mechanisms, both in the setting of this competition and in the more general case.

251B
Mobile Opportunistic Commerce: Mechanisms, Architecture, and Application
Ece Kamar, Eric Horvitz, Chris Meek
We present mechanisms, architectures, and an implementation addressing challenges with mobile opportunistic commerce centering on markets and mechanisms that support the procurement of goods and services in mobile settings. Our efforts seek to extend

core concepts from research in electronic commerce to interactions between mobile buyers and brick and mortar businesses that have geographically situated retail offices. We focus on efficient mechanisms, infrastructure, and automation that can enable sellers and buyers to take joint advantage of the relationship of the locations of retail offices to the routes of mobile buyers who may have another primary destination. The methods promote automated vigilance about opportunities to buy and sell, and to support negotiations on the joint value to buyers and sellers including buyers' costs of divergence from their original paths to acquire services and commodities. We extend prior work on auction mechanisms to personal procurement settings by analyzing the dynamics of the cost to buyers based on preexisting plans, location, and overall context. We present mechanisms for auctions in single item, combinatorial, and multiattribute settings that take into consideration personal inconvenience costs within time-sensitive dynamic markets and challenges with privacy and fairness.

253A

Selecting Strategies Using Empirical Game Models: An Experimental Analysis of Meta-Strategies

Christopher Kiekintveld, Michael Wellman

In many complex multi-agent domains it is impractical to compute exact analytic solutions. An alternate means of analysis applies computational tools to derive and analyze empirical game models. These models are noisy approximations, which raises questions about how to account for uncertainty when analyzing the model. We develop a novel experimental framework and apply it to benchmark meta-strategies – general algorithms for selecting strategies based on empirical game models.

We demonstrate that modeling noise is important; a naïve approach that disregards noise and plays according to Nash equilibrium yields poor choices. We introduce three parameterized algorithms that factor noise into the analysis by predicting distributions of opponent play. As observation noise increases, rational players generally make less specific outcome predictions. Our comparison of the algorithms identifies logit equilibrium as the best method for making these predictions. Logit equilibrium incorporates a form of noisy decision-making by players. Our evidence shows that this is a robust method for approximating the effects of uncertainty in many contexts. This result has practical relevance for guiding analysis of empirical game models. It also offers an intriguing rationale for behavioral findings that logit equilibrium is a better predictor of human behavior than Nash equilibrium.

Friday, May 16th, 13:40-14:40 – Session 8

Multi-Robotics III

Chair: Radhika Nagpal

****38B****

Sensing-based Shape Formation on Modular Multi-Robot Systems: A Theoretical Study

Chih-Han Yu, Radhika Nagpal

This paper presents a theoretical study of decentralized control for sensing-based shape formation on modular multi-robot systems, where the desired shape is specified in terms of local sensor constraints between neighboring robot agents. We show that this problem can be formulated more generally as "distributed constraint-maintenance" on a networked multi-agent system. It is strongly related to a class of multi-agent algorithms called "distributed consensus", which includes several bio-inspired algorithms such as flocking and firefly synchronization. By exploiting this connection, we can theoretically analyze several important aspects of the decentralized shape formation algorithm and generalize it to more complex multi-agent scenarios. We show that the convergence time depends on (a) the number of robot agents and agent connection topology, (b) the complexity of the user-specified goal, and (c) the initial state of the robots. Using these results, we can provide precise statements on how the approach scales, and how quickly the system can adapt to perturbations. These results provide a deeper understanding of the contrast between centralized and decentralized multi-agent algorithms.

41B

PNP: A Formal Model for Representation and Execution of Multi-Robot Plans.

Vittorio Ziparo, Luca Iocchi, Daniele Nardi, Pier Francesco Palamara, Hugo Costelha

High level programming of mobile robots is very important for developing complex and reliable robotic applications. High level programming is usually performed by defining plans, i.e. program structures describing action execution control. To develop complex applications plans should be able to represent many features such as sensing, loops, concurrency, non-instantaneous actions, action failures, and action synchronization in a multi-agent context.

The aim of this paper is to describe a novel representation framework for high level robot and multi-robot programming, called Petri Net Plans (PNP), that allows for representing all the action features that are needed for describing complex plans in dynamic environments. PNPs are based on Petri nets, a graphical modeling language for dynamic systems. We provide a sound and complete execution algorithm for PNPs based on the semantics of Petri nets. Moreover, we show that multi-robot PNPs allow for a sound and complete distributed execution algorithm, given that a reliable communication channel is provided.

PNPs have been used for describing effective plans for actual robotic agents which inhabit dynamic, partially observable and unpredictable environments, and experimented in different application scenarios, including robotic soccer, rescue competitions and a complex multi-robot foraging task.

Agent-based System Development IV

Chair: Michael Winikoff

58A

Towards verifying compliance in agent-based web service compositions

Alessio Lomuscio, Hongyang Qu, Monika Solanki

We explore the problem of specification and verification of compliance in agent based Web service compositions. We use the formalism of temporal-epistemic logic suitably extended to deal with compliance/violations of contracts. We illustrate these concepts using a motivating example where the behaviours of participating agents are governed by contracts. The composition is specified in OWL-S and mapped to our chosen formalism. Finally we use an existing symbolic model checker to verify the example specification whose state space is approximately 2^{21} and discuss experimental results.

61A

Semantic Matchmaking of Web Services using Model Checking

Akin Gunay, Pinar Yolum

Service matchmaking is the process of finding suitable services given by the providers for the service requests of consumers. Previous approaches to service matchmaking is mostly based on matching the input-output parameters of service requests and service provisions. However, such approaches do not capture the semantics of the services and hence cannot match requests to services precisely. This paper we propose an agent-based approach for matchmaking that is based on capturing the semantics of services and requests formally through temporal logic. Requests are represented as set of properties and compared to the service representations using model checking, yielding results on whether a service can satisfy a request or not. By help of domain ontologies, our approach also supports flexible matching, where partially matching services are identified. We provide a general framework, where our approach can work with other existing matchmaking approaches and is integrated with current efforts such as OWL-S and SWRL. We conduct a detailed case study to evaluate our approach with respect to other matchmaking approaches.

Economic Paradigms VIII

Chair: Felix Brandt

227B

Distributed Multiagent Resource Allocation in Diminishing Marginal Return Domains

Yoram Bachrach, Jeffrey Rosenschein

We consider a multiagent resource allocation domain where the marginal production of each resource is diminishing. A set of identical, self-interested agents requires access to sharable resources in the domain. We present a distributed and random allocation procedure, and demonstrate that the allocation converges to the optimal in terms of utilitarian social welfare. The procedure is based on direct interaction among the agents and resource owners (without the use of a central authority).

We then consider potential strategic behavior of the self-interested agents and resource owners, and show that when both act rationally and the domain is highly competitive for the resource owners, the convergence result still holds. The optimal allocation is arrived at quickly; given a setting with k resources and n agents, we demonstrate that the expected number of timesteps to convergence is O(k ln n), even in the worst case, where the optimal allocation is extremely unbalanced.

Our allocation procedure has advantages over a mechanism design approach based on Vickrey-Clarke-Groves (VCG) mechanisms: it does not require the existence of a central trusted authority, and it fully distributes the utility obtained by the agents and resource owners i.e., it is strongly budget-balanced).

235A

Trajectories of Goods in Distributed Allocation

Yann Chevaleyre, Ulle Endriss, Nicolas Maudet

Distributed allocation mechanisms rely on agents' autonomous (and supposedly rational) behaviour: states evolve as a result of agents' contracting deals and exchanging resources. It is no surprise that restrictions on agents' potential deals restrict the reachability of some desirable states, for instance states where goods are efficiently allocated. In particular *topological* restrictions make any attempt to guarantee asymptotic convergence to an optimal allocation vacuous in most cases. In this paper, we concentrate instead on the *dynamics* of such systems; more precisely we study the trajectories of goods in such iterative reallocative processes. Our first contribution is to propose an upper bound on the length of goods' trajectories, when agents' utility functions are modular. The second innovative aspect of the paper is then to discuss how this affects, on average, the quality of the states that are reached. Finally, a preliminary study of the non-modular case is proposed, examining how synergetic effects between items can affect their trajectories.

Short Papers

Multi-Robotics Track

45C

On Reduced Time Fault Tolerant Paths for Multiple UAV Covering a Hostile Terrain

Rahul Sawhney, Madhava Krishna, Srinathan Kannan, Mahesh Mohan

We present a method for finding reduced time coverage paths of multiple UAVs (Unmanned Air Vehicles) monitoring a 3D terrain represented as height fields. A novel metric based on per time visibility is used that couples visibility gained at a terrain point with the time spent to reach the point. This coupled metric is utilized to form reduced time paths by maximizing the visibility gained per unit time at every step. We compare the results of this approach with an approach that covers the terrain based on a per distance visibility metric, which reduces the sum, over distances covered by each UAV path. The comparisons show that the current method gives substantially time reduced paths albeit with an expected increase in sum over distances of UAV paths. We also show that time taken to cover the terrain based on the current metric is far less than prevalent methods that try to decompose the terrain based on visibility followed by time or time followed by visibility in a decoupled fashion. The method is further extended to provide for fault tolerance on a hostile terrain. Each terrain point is guaranteed to be seen by at-least one UAV that has not been damaged due to any calamity, shot or otherwise.

46A

Protoswarm: A Language for Programming Multi-Robot Systems Using the Amorphous Medium Abstraction

Jonathan Bachrach, James McLurkin, Anthony Grue

Multi-robot systems are becoming increasingly prevalent, but programmability is a major barrier to their deployment. Present systems force programmers to think in terms of individual agents. Application code becomes entangled with details of coordination and robustness and often does not compose well or translate to other domains. We offer an alternate approach whereby the programmer controls a single virtual *spatial computer* which fills the environment space. The computations on this spatial computer are actually performed by a large number of locally-interacting individual agents. This abstracts the actual computational hardware behind the spatial computer interface, and allows the programmer to

focus on a single model of global computation. We achieve this abstraction with two components: a language that embodies continuous space and time semantics and a runtime library that implements these semantics approximately. We demonstrate the efficacy of our approach with multi-agent algorithms in both simulation and on a group of 40 robots.

47B

A Scalable and Distributed Model for Self-Organization and Self-Healing

Michael Rubenstein, Wei-Min Shen

As the ability to produce a large number of small, simple robotic agents improves, it becomes essential to control the behavior of these agents in such a way that the sum of their actions gives rise to the desired overall result. These agents are modeled as homogenous, distributed robots, with only one simple sensor, with a shorter than global range. Our simple agents are tasked to form and hold a desired swarm shape, independent of the total number of agents. If this shape is damaged by the removal of some of the agents, the remaining agents will recover the shape, but with a smaller scale. These shapes can also have a pattern such as a picture or drawing displayed on them by controlling the individual robots color, symbolically representing the differentiation of agents within the swarm. This pattern will resize to fit the existing swarm. With the ability to synchronize in time, the swarm gains the ability to change the pattern displayed, resulting in a moving image.

48C

Teaching Multi-Robot Coordination using Demonstration of Communication and State Sharing

Sonia Chernova, Manuela Veloso

Solutions to complex tasks often require the cooperation of multiple robots, however, developing multi-robot policies can present many challenges. In this work, we introduce teaching by demonstration in the context of multi-robot tasks, enabling a single teacher to instruct multiple robots to work together through a demonstration of the desired behavior. Within this framework, we contribute two approaches for teaching coordination based on different communication and information sharing strategies. To enable the teacher to divide attention between multiple robots, each robot uses a confidence-based algorithm that allows it to regulate its autonomy and determine the need for demonstration. Evaluation is performed using two Sony QRIO robots learning a real-world collaborative ball sorting task.

49A

Modelling, Analysis and Execution of Multi-Robot Tasks using Petri Nets

Hugo Costelha, Pedro Lima

This paper introduces Petri net (PN) based models of cooperative robotic tasks, namely those involving the coordination of two or more robots, thus requiring the exchange of synchronisation messages, either using explicit (e.g., wireless) or implicit (e.g., vision-based observation of teammates) communication. In the models, PN places represent primitive actions, subtasks and predicates set by sensor readings and communicated messages. Events are associated to PN transitions. The PN models can be used for task planning, plan execution and plan analysis. Different PN views enable the analysis of different properties. In this work we focus on plan analysis, namely on properties such as boundedness and liveness, corresponding to checking if resources usage is stable and plans have no deadlocks, as well as on stochastic performance, concerning the plan success probability. One novel feature of our work is that the analysis consists of composing several small action PN models with environment PN models, leading to a closed loop robot team/environment analysis methodology. Examples of application to simulated robotic soccer scenarios are presented.

50B

Coalition game-based distributed coverage of unknown environments by robot swarms

Ke Cheng, Prithviraj Dasgupta

We consider the problem of distributed exploration or coverage of an unknown environment by a swarm of mobile mini-robots that have limited memory, computation and communication capabilities. We describe a novel mechanism of distributed coverage of an unknown environment by swarmed robots that can dynamically merge and split into structured teams or exchange team members to improve the efficiency of solving the coverage problem. Our mechanism combines the technique of swarm-based flocking with coalition games to enable robots dynamically select utility maximizing teams that move in formation.

51C

Decentralized Coordination of Automated Guided Vehicles

David Herrero-Pérez, Humberto Martínez-Barberá

This paper approaches the issue of coordination of highly autonomous Automated Guided Vehicles (AGVs) working on an automated factory. These vehicles are used for goods delivery tasks between different points of the production system. The coordination is based on a decentralized architecture where each vehicle broadcasts the information about its state in the working environment, and by combining all these states in a local way, each AGV decides which action to take. The heuristic that allows the decentralized traffic control is based on a priority system, based on the current task, and a set of dangerous zones which are defined to avoid possible deadlocks, where mutual exclusion should be ensured. The process is somehow similar to that used by humans when circulating in cars: a set of rules and a set of signals/places. The interaction of many vehicles working on the same area under different collision conditions has been tested in a real industrial warehouse environment.

52A

Robust Team Play in Highly Uncertain Environments

Henry Work, Eric Chown, Tucker Hermans, Jesse Butterfield

Effective teamwork in highly dynamic environments requires a delicate balance between giving agents the autonomy to act and react on their own and restricting that autonomy so that the agents do not work at cross purposes. In this article we describe the problems involved in coordinating behavior based upon a highly dynamic object, a soccer ball, for agents with sensing and communications limitations. We then present a system for coping with these problems and examine its success in light of its performance at the RoboCup 2007 championship.

53B

A Coordination Mechanism for Swarm Navigation: Experiments and Analysis

Leandro Marcolino, Luiz Chaimowicz

We present an algorithm that allows swarms of robots to navigate in environments containing unknown obstacles, moving towards and spreading along 2D shapes given by implicit functions. Basically, a gradient descent approach augmented with local obstacle avoidance is used to control the swarm. To deal with local minima regions, we use a coordination mechanism that reallocates some robots as "rescuers" and sends them to help other robots that may be trapped. The main objective of this paper is to analyze the performance of this coordination algorithm in terms of its completion rate and communication requirements as the number of robots increases. For this, a series of simulations are presented and discussed.

54C

openSDK - An Open-source Implementation of OPEN-R

Nuno Lopes, Pedro Lima

This paper describes openSDK, an open-source implementation of Sony's AIBO development kit (OPEN-R). openSDK is capable of running unmodified AIBO programs (only a recompilation is necessary) on a standard computer, us- ing a simulator at full frame rate (currently only USARSim is supported) or on a different robotic hardware platform. openSDK also offers standard debugging facilities for AIBO programs.

55B

Multi-robot Markov Random Fields

Jesse Butterfield, Odest Jenkins, Brian Gerkey

We propose using the Markov Random Field (MRF) as a probabilistic mathematical model for unifying approaches to multi-robot coordination, or, more specifically, distributed action selection. The MRF model is suited towards domains where the joint probability over latent (action) and observed (perceived) variables can be factored into pairwise interactions between these variables. Specifically, these interactions occur through functions which evaluate "local evidence" between an observed and latent variable and "compatibility" between a pair of latent variables. For multi-robot coordination, we cast local evidence functions as the computation for an individual robot's action selection from its local observations and compatibility as the dependence in action selection between a pair of robots. We describe how existing methods for multi-robot coordination (or at least a non-exhaustive subset) fit within an MRF-based model and how they conceptually unify. Further, we offer belief propagation on a multi-robot MRF as a new alternative to existing distributed allocation algorithms.

Virtual Agents Track

14B

Creating Crowd Variation with the OCEAN Personality Model

Funda Durupinar, Jan Allbeck, Nuria Pelechano, Norman Badler

Most current crowd simulators animate homogeneous crowds, but include underlying parameters that can be tuned to create variations within the crowd. These parameters, however, are specific to the crowd models and may be difficult for an animator or naïve user to use. We propose mapping these parameters to personality traits. In this paper, we extend the HiDAC (High-Density Autonomous Crowds) system by providing each agent with a personality model in order to examine how the emergent behavior of the crowd is affected. We use the OCEAN personality model as a basis for agent psychology. To each personality trait we associate nominal behaviors; thus, specifying personality for an agent and leading to an automation of the low-level parameter tuning process. We describe a plausible mapping from personality traits to existing behavior types and analyze the overall emergent crowd behaviors.

17B

Influence of Social Relationships on Multiagent Persuasion

Katsunori Kadowaki, Kazuki Kobayashi, Yasuhiko Kitamura

Life-like agents have potential to make e-shopping sites on the Web more attractive and persuasive and we have interest in how multiple life-like agents should behave as a team to persuade customers. To know how the social relation among two agents and a human user affects the performance of persuasion from a viewpoint of the balance theory, we develop a multi-agent persuasion system. In the system, the agents construct a social relation to the user, and then they try to persuade him/her to select items that they recommend. The evaluation result shows that the performance of agents in the balanced relation exceeds that in the imbalanced one.

19A

What Should the Agent Know? The Challenge of Capturing Human Knowledge

Emma Norling

Reports of applications that include agent-based models of human behaviour tend to focus on the applications themselves and the success of the modelling exercise. They give little (if any) information on the process used to construct the models. Those who attempt to construct models of human behaviour quickly realise that this is a non-trivial task. This paper presents a methodological approach to eliciting knowledge for BDI-based models of human behaviour.

20B

iCat: An Affective Game Buddy Based on Anticipatory Mechanisms

Iolanda Leite, Carlos Martinho, André Pereira, Ana Paiva

In this paper, we study the role of emotions and expressive behaviour in socially interactive characters employed in educational games. More specifically, on how we can use such emotional behaviour to help users to better understand the game state. An emotion model for these characters, which is mainly influenced by the current state of the game and is based on the emotivector anticipatory mechanism, was developed. We implemented the model in a social robot named iCat, using chess as the game scenario. The results of a preliminary evaluation suggested that the emotional behaviour embedded in the character indeed helped the users to have a better perception of the game.

21C

If I were you - Double appraisal in affective agents

Ruth Aylett, Sandy Louchart

This paper reports on the implementation and evaluation of a Simulation Theory (ST) approach to the Theory of Mind in intelligent graphical agents driven by an

affective agent architecture FatiMA. The cognitive appraisal mechanism already present in FatiMA is adapted to produce a second appraisal cycle, a double appraisal, in order to evaluate the emotional impact of possible actions. The action with the greatest emotional impact is selected as a means of producing more interesting dramatic actions. A variant in which the actual minds of characters present are used is also implemented and evaluated. Results show that these mechanisms do produce more interesting stories.

22A

User's Gestural Exploration of Different Virtual Agents' Expressive Profiles

Matthieu Courgeon, Jean-Claude Martin, Christian Jacquemin

Designing affective user interfaces involving expressive virtual characters raises several research questions. From a computer science point of view, the system should be able to display facial expressions of emotion as dynamic realtime reactions to user's inputs. From a perception point of view, designers need to know how the user will perceive the dynamics of these facial expressions in relation with his own input. We aim at studying if users are able to perceive different expressive profiles of virtual characters from a low-level interaction. This paper describes our platform that enables a virtual character to display facial expressions of emotions as realtime continuous reactions to users' gesture input. We detail the graphical rendering components and explain the techniques underlying the computation of intermediate facial expressions of emotion, and their control in the 3D P.A.D. space using gesture input via a joystick. Preliminary results of a perception study are presented and show the potential of such an approach for studying the dynamics of the perception of emotional expressions during interaction with virtual characters endowed with different expressive profiles.

23B

Negotiating Task Interruptions with Virtual Agents for Health Behavior Change

Timothy Bickmore, Daniel Mauer, Francisco Crespo, Thomas Brown

Virtual health counseling agents on mobile devices need to be able to interrupt their users when it is time for them to engage in healthy behaviors, such as scheduled medication taking or exercise. However, these real-time reminders often represent task interruptions for individuals who are engaged in work activities. This paper presents the results of a study which compares four strategies used by a virtual agent on a PDA for interrupting users at work to perform a healthy behavior. We find that, among several interruption coordination strategies previously explored in the HCI literature, empathic interruptions are superior overall in gaining both short-term compliance and self-reported desire to continue working with the agent.

24C

The Senior Companion Multiagent Dialogue System

Hugo Pinto, Yorick Wilks, Roberta Catizone, Alexiei Dingli

This article presents a multiagent dialogue system. We show how a collection of relatively simple agents is able to treat complex dialogue phenomena and deal sucessfully with different deployment configurations. Next, we compare this multi-agent realization to a previous monolithic implementation where the treatment of dialogue phenomena was explicitly hard-coded, along the dimensions of configurability, robustness and emergent behavior. Finally, we situate our work among other multiagent dialogue systems.

25A

Dancing the Night Away — Controlling a Virtual Karaoke Dancer by Multimodal Expressive Cues

Matthias Rehm, Thurid Vogt, Nikolaus Bee, Michael Wissner

In this article, we propose an approach of nonverbal interaction with virtual agents to control agents' behavioral expressivity by extracting and combining acoustic and gestural features. The goal for this approach is twofold, (i) expressing individual features like situated arousal and personal style and (ii) transmitting this information in an immersive 3D environment by suitable means.

26B

iCat, the Chess Player: The influence of embodiment in the enjoyment of a game

André Pereira, Carlos Martinho, Iolanda Leite, Ana Paiva

This paper presents an experiment that evaluates and compares the user enjoyment when playing a game of chess in two situations: against a physically embodied robotic agent and against a virtually embodied agent, displayed on screen. The results of the study suggest that embodiment has implications on user enjoyment, as the experience against a robotic agent was classified as more enjoyable than against a virtually embodied agent.

27C

Learning to Interact: Connecting Perception with Action in Virtual Environments

Pedro Sequeira, Ana Paiva

Modeling synthetic characters which interact with objects in dynamic virtual worlds is important when we want the agents to act in an autonomous and non-preplanned way. Such interactions with objects would allow the synthetic characters to behave in a more believable way. Once objects offer innumerous uses, it is essential that the agent is able to acquire the necessary knowledge to identify action possibilities in the objects while interacting with them. We propose a conceptual framework that allows the agents to identify possible interactions with objects based in past experiences with other objects. Starting from sensory patterns collected during interactions with objects, the agent is able to acquire conceptual knowledge about regularities of the world, its internal states and its own actions. The presented work also proposes that such acquired knowledge may be used by the agent in order to satisfy its needs and goals by interacting with objects. Preliminary tests were made and it is possible to state that our agents are able to acquire valid conceptual knowledge about the regularities in the environment and its objects, its own actions and causal relations between them.

28A

The Intensity of Perceived Emotions in 3D Virtual Humans

Ahmad Shaarani, Daniela Romano

Synthetically generated 3D humans often fail to express a full range of emotions or present different levels of the same type of emotion. Transcending the facial expression, what should a happy synthetically generated human look like? What about a slightly happy or

ecstatically happy? This paper reports a study aimed at identifying the appropriate bodily expressions for various emotions in 3D human-like figures at varying emotional strength. Thirty-six volunteers were asked to discriminate and categorize thirty cards with static poses of 3D human-like characters into the Ekman's six basic categories of emotions. This is to judge the compatibility of each posture in relation to each category and to rate their level of emotion within the group.

29B

Conviviality Masks in Multiagent Systems

Patrice Caire, Serena Villata, Guido Boella, Leon van der Torre

In this paper we study tools for conviviality to develop user-friendly multiagent systems. First, we show how to use the social-cognitive concept of conviviality in multiagent system technology by relating it to agent power and social dependence networks. Second, we define conviviality masks as transformations of social dependencies by hiding power relations and social structures to facilitate social interactions. Third, we introduce dynamic dependence networks to model the creation of conviviality using conviviality masks. We illustrate the use of conviviality masks with a multiagent teleconferencing application for virtual worlds.

30C

So tell me what happened: turning agent-based interactive drama into comics

Tiago Alves, Ana Rita Simões, Rui Figueiredo, Marco Vala, Ana Paiva, Ruth Aylett

As virtual characters become more autonomous, their use in interactive drama is growing. By creating interesting and well authored personalities, these characters are able to interact with each other and create appealing non-scripted stories. As with any other story, we may want to re-tell it or create a summary of the essential parts. Our goal was to create comic-like summaries of these stories. This paper presents a system that analyses story logs, looks at the characters' emotional information to understand their actions and their importance in the story, selects the most important events and creates comic strips. Forty users evaluated the system and the results show that the summaries driven by the characters' emotional information are effective when compared with showing the stories as they unfold.

31A

Methods for Complex Single-Mind Architecture Designs (Short Paper)

Kristinn Thorisson, Gudny Ragna Jonsdottir, Eric Nivel

The implementation of software systems with large numbers of heterogeneous components calls for a powerful design methodology. Although several such methodologies have been proposed, many lack application to construction of single-mind systems. We have employed the Constructionist Design Methodology (CDM) in building several such systems, including an autonomous radio show host. Proposing modules communicating through messages via blackboards as key building blocks for interactive intelligences, the methodology has been of considerable help in the early stages of designing several large architectures. This paper describes efforts to extend the CDM with more detailed support for the modularization process. We describe our use of a combination of abstraction and finite state machines in modularizing the realtime turntaking system of the radio show host. Our experience shows considerable benefits and added flexibility in the creation of large architectures when using the new modularization principles.

32B

The Intermediary Agent's Brain: Supporting Learning to Collaborate at the Inter-Personal Level

Juan Martinez-Miranda, Bernhard Jung, Sabine Payr, Paolo Petta

We discuss the design of the Intermediary Agent's brain, the control module of an embodied conversational virtual peer in a simulation game aimed at providing learning experiences regarding the dynamics of collaboration at the inter-personal (IP) level. We derive the overall aims of the game from theoretical foundations in collaboration theory and pedagogical theory and related requirements for the virtual peer; present the overall modular design of the system; and then detail the design perspectives and the interplay of the related operationalised concepts leading to the control architecture of the Intermediary Agent, that is realised as a simple cognitive appraisal process driven by direct and indirect effects of the mission-oriented and social interactions of players and agent on the agent's level of trust in its human peers. We conclude with coverage of related work and insights from first deployment experiences.

33C

A 'Companion' ECA with Planning and Activity Modelling

Marc Cavazza, Cameron Smith, Daniel Charlton, Li Zhang, Jaakko Hakulinen, Markku Turunen

In this paper, we describe the development of an Embodied Conversational Agent (ECA) implementing the concept of a companion, i.e. an agent supporting the persistent representation of user activities and dialogue-based communication with the user. This first experiment implements an Health and Fitness companion aimed at promoting a more healthy lifestyle. The system operates by generating an ideal plan of daily activities from background knowledge and dialogue interaction with the user. This plan then becomes an activity model, which will later be instantiated by reports from the user and analysed by the agent from the perspective of initial objectives. At various stages of the day, the plan can still be adapted through further dialogue. The agent is embodied using a wireless rabbit (Nabaztag™) device situated in the user's home. After describing the planning component, based on Hierarchical Task Networks and the spoken dialogue system, we present working examples from the system illustrating its behaviour through various phases of user activity generation, updating and re-planning.

34A

Emotional Reading of Medical Texts Using Conversational Agents

Gersende Georg, Catherine Pelachaud, Marc Cavazza

In this paper, we present a prototype that helps visualizing the relative importance of sentences extracted from medical texts using Embodied Conversational Agents (ECA). We propose to map rhetorical structures automatically recognized in the documents onto a set of communicative acts controlling the expression of an ECA. As a consequence, the ECA will dramatize a sentence to reflect its perceived importance and rhetorical strength (advice, requirement, open proposal, etc). This prototype is constituted of three sub-systems: i) G-DEE, a text analysis module ii) a mapping module which

converts rhetorical structures produced by the text analysis module into communicative functions driving the ECA animation and iii) an ECA system. By bringing the text to life, this system could help their authors (in our application, expert physicians) to reflect on the potential impact of the writing style they have adopted. The use of ECA re-introduces an affective element which cannot easily be captured by other methods for analyzing document's style.

35B

Individual Differences in Expressive Response: A Challenge for ECA Design

Ning Wang, Stacy Marsella, Tim Hawkins

To create realistic and expressive virtual humans, we need to develop better models of the processes and dynamics of human emotions and expressions. A first step in this effort is to develop means to systematically induce and capture realistic expressions in real humans. We conducted a series of studies on human emotions and facial expression using the Emotion Evoking Game (EVG) and a high-speed video camera. EVG allows researchers to systematically explore factors that elicit emotions. In this paper, we discuss a detailed analysis of facial expressions in response to a surprise situation. We found individual differences in whether surprise was evoked and how it was expressed. We also provide details on the rich dynamics of facial expressions.

36C

Another Look at Search-Based Drama Management

Michael Mateas

A drama manager (DM) is a system that monitors an interactive experience, such as a computer game, and intervenes to keep the global experience in line with the author's goals without decreasing a player's interactive agency. In declarative optimization-based drama management (DODM), an author declaratively specifies desired properties of the experience; the DM intervenes in a way that optimizes the specified metric. The initial DODM approach used online search to optimize an experience-quality function. Later work questioned both online search as a technical approach and the experience-quality optimization framework. Recent work on targeted trajectory distribution Markov decision processes (TTD-MDPs) replaced the experience-quality metric with a metric and associated algorithm based on targeting experience distributions. We show that, though apparently quite different on the surface, the original optimization formulation and TTD-MDPs are actually variants of the same underlying search algorithm, and that offline cached search, as is done by the TTD-MDP algorithm, allows the original search-based systems to achieve similar results to TTD-MDPs. Furthermore, we argue that the original idea of optimizing an experience-quality function does not destroy interactive agency, as had previously been argued, and that in fact it can capture that goal directly.

Agent-Based System Development

64A

On the Importance of Migration for Fairness in Online Grid Markets

Lior Amar, Ahuva Mu'alem, Jochen Stößer

Computational grids offer users a simple access to tremendous computer resources for solving large scale computing problems. Traditional performance analysis of scheduling algorithms considers overall system performance while fairness analysis focuses on the individual performance each user receives. Until recently, only few grids and cluster systems provided preemptive migration, which is the ability of dynamically moving computational tasks across machines during runtime. The emergent technology of virtualization provides off-the-shelf support for migration, thus making the use of this feature more accessible (even across different OS's).

In this paper, we study the close relation between migration and fairness. We present fairness and quality of service properties for economic online scheduling algorithms. Under mild assumptions we show that it is impossible to achieve these properties without the use of migration. On the other hand, if zero cost migration is used, then these properties can be satisfied.

65B

A Flexible Framework for Verifiable Agent Programming

Louise Dennis, Berndt Farwer, Rafael Bordini, Michael Fisher

There is an increasing number of agent-oriented programming languages that have working interpreters and platforms, with significant progress in the quality of such platforms over the last few years. With these platforms becoming more popular, and multi-agent systems being increasingly used for safety-critical applications, the need for verification techniques that apply to systems written in such languages is proportionally intensified. Building on our previous work on model checking for a particular agent-oriented programming language, we have developed a new approach whereby model checking techniques can be used directly on a variety of such languages. The approach also supports the verification of multi-agent systems where individual agents have been programmed in different agent languages.

66C

Structure in Threes: Modelling Organization-Oriented Software Architectures Built Upon Multi-Agent Systems

Matthias Wester-Ebbinghaus, Daniel Moldt

Software systems are subject to increasing complexity and in need of efficient structuring. Multi-agent system research has come up with approaches for an organization-oriented comprehension of software systems. However, when it comes to the collective level of organizational analysis, multi-agent system technology lacks clear development concepts. To overcome this problem while preserving the earnings of the agent-oriented approach, this paper propagates a shift in perspective from the individual agent to the organization as the core metaphor of software engineering targeting at very large systems. According to different levels of analysis drawn from organization theory, different types of organizational units are incorporated into a reference architecture for organization-oriented software systems.

67A

Engineering Large-scale Distributed Auctions

Peter Gradwell, Michel Oey, Reinier Timmer, Julian Padget, Frances Brazier

The functional characteristics of market-based solutions are typically best observed through the medium of simulation, data-gathering and subsequent visualization. We previously developed a simulation of multiple distributed auctions to handle resource allocation (in fact,

bundles of unspecified goods) and in this paper we want to deploy an equivalent system as a distributed application. There are two notable problems with the simulation-first, application-second approach: (i) the simulation cannot reasonably take account of network effects, and (ii) how to recreate in a distributed application the characteristics demonstrated by the mechanism in the simulation. We describe: (i) the refactorings employed in the process of transforming a uni-processor lock-step simulation into a multi-processor asynchronous system, (ii) some preliminary performance indicators, and (iii) some reflections on our experience which may be useful in building MAS in general.

68B
Ontology-based Test Generation for Multi Agent Systems
Cu Nguyen, Anna Perini, Paolo Tonella
Software agents are a promising technology for today's complex, distributed systems. Hence, methodologies and techniques that address testing and reliability of multi agent systems are increasingly demanded, in particular to support automated test case generation and execution. Agent testing is aimed at exercising and stressing the agents under test, so as to bring them to situations where misbehaviors are revealed. In this paper, we take advantage of agent interaction ontologies that define content semantic of agent interactions to: (i) generate test inputs; (ii) guide the exploration of the input space during generation; and, (iii) verify messages exchanged among agents with respect to the defined interaction ontology. We integrated the proposed approach into a novel testing framework, called eCAT, which can generate and evolve test cases automatically, and run them continuously. We used two BDI agent applications as case studies to illustrate the performance of the framework as well as its capability to reveal faults.

Agent and Multi-Agent Learning

208A
Non-linear Dynamics in Multiagent Reinforcement Learning Algorithms
Sherief Abdallah, Victor Lesser
Several multiagent reinforcement learning (MARL) algorithms have been proposed to optimize agents\' decisions. Only a subset of these MARL algorithms both do not require agents to know the underlying environment and can learn a stochastic policy (a policy that chooses actions according to a probability distribution). Weighted Policy Learner (WPL) is a MARL algorithm that belongs to this subset and was shown, experimentally in previous work, to converge and outperform previous MARL algorithms belonging to the same subset.
An important aspect of understanding the behavior of a MARL algorithm is analyzing the dynamics of the algorithm: how policies of multiple learning agents evolve over time while interacting with one another. Such an analysis reveals whether agents using a particular MARL algorithm will eventually converge. The dynamics analysis also points out features of the MARL algorithm that are exhibited during the convergence period.
The main contribution of this paper is analyzing the dynamics of WPL and showing the effect of its non-linear nature, as opposed to previous MARL algorithms that had linear dynamics. First, we represent the WPL algorithm as a set of differential equations. We then solve the equations and show that it is consistent with experimental results reported in previous work. We finally compare the dynamics of WPL with earlier MARL algorithms and discuss the interesting differences and similarities we have discovered.

209B
Expediting RL by Using Graphical Structures
Peng Dai, Alexander Strehl, Judy Goldsmith
Reinforcement learning (RL) is a means of learning an optimal policy for Markov decision problems without knowing the underlying model a priori. However, RL algorithms are usually magnitudes slower than planning algorithms since they must base their planning on learning experiences. Recent literature shows that MDP planning can be significantly expedited by making use of the graphical structure of the MDP. In this paper, we propose extensions to two popular RL algorithms, Q-learning and RMax, by learning the graphical structure of problems and using the structure to improve overall learning speed. Use of the graphical structure of the underlying MDP can greatly improve the speed of planning algorithms, if the underlying MDP has a nontrivial topological structure. However, our experiments show that use of the apparent topological structure of an MDP speeds up machine learning, even if the MDP is simply connected.

210C
Transfer of Task Representation in Reinforcement Learning using Policy-based Proto-value Functions
Eliseo Ferrante, Alessandro Lazaric, Marcello Restelli
Reinforcement Learning research is traditionally devoted to solve single-task problems. This means that, anytime a new task is faced, learning must start from scratch. Recently, several studies have addressed the issues of reusing the knowledge acquired in solving previous related tasks by transferring information about policies and value functions. In this paper we analyze the use of proto-value functions under the transfer learning perspective. Proto-value functions are effective basis functions for the approximation of value functions defined over the graph obtained using the random walk on the environment. The building of this graph is a key aspect for transfer learning; in particular, when we want to transfer knowledge between tasks that have different reward function and domains, it is important that the graph captures information from both these elements. For this reason, we introduce policy-based proto-value functions, which can be obtained by considering the graph generated by the walk on the environment guided by the optimal policy of a task at hand. To compare the effectiveness of the two kinds of proto-value functions, we show experimental results on a simple grid-world environment where we set up both goal-transfer and domain-transfer problems.

211A
Reinforcement Learning for DEC-MDPs with Changing Action Sets and Partially Ordered Dependencies
Thomas Gabel, Martin Riedmiller
Decentralized Markov decision processes are frequently used to model cooperative multi-agent systems. In this paper, we identify a subclass of general DEC-MDPs that features regularities in the way agents interact with one another. This class is of high relevance for many real-world applications and features provably reduced

complexity (NP-complete) compared to the general problem (NEXP-complete). Since optimally solving larger-sized NP-hard problems is intractable, we keep the learning as much decentralized as possible and use multi-agent reinforcement learning to improve the agents' behavior online. Further, we suggest a restricted message passing scheme that notifies other agents about forthcoming effects on their state transitions and that allows the agents to acquire approximate joint policies of high quality.

212B

Using Adaptive Consultation of Experts to Improve Convergence Rates in Multiagent Learning

Greg Hines, Kate Larson

We present a regret-based multiagent learning algorithm which is provably guaranteed to converge (during self-play) to the set of Nash equilibrium in a wide class of games. Our algorithm, FRAME, consults experts in order to obtain strategy suggestions for agents. If the experts provide effective advice for the agent, then the learning process will quickly reach a desired outcome. If, however, the experts do not provide good advice, then the agents using our algorithm are still protected. We further expand our algorithm so that agents learn, not only how to play against the other agents in the environment, but also which experts are providing the most effective advice for the situation at hand.

213C

A New Perpective to the Keepaway Soccer: The Takers

Atil Iscen, Umut Erogul

Keepaway is a sub-problem of RoboCup Soccer Simulator in which 'the keepers' try to maintain the possession of the ball, while 'the takers' try to steal the ball or force it out of bounds. By using Reinforcement Learning as a learning method, a lot of research has been done in this domain. In these works, there has been a remarkable success for the intelligent keepers part, however most of these keepers are trained and tested against simple hand-coded takers. We tried to address this part of the problem by using Sarsa(λ) as Reinforcement Learning method with linear tile-coding as function approximation and used two different state spaces that we specially designed for the takers. As the results of the experiments confirm; we created a better trainer and tester for the keepers, and also outperformed the hand-coded taker. Also when designing the new state space, we noticed that smaller state spaces can also be successful for this part of the problem.

214A

On the Usefulness of Opponent Modeling: the Kuhn Poker case study

Alessandro Lazaric, Mario Quaresimale, Marcello Restelli

The application of reinforcement learning algorithms to Partially Observable Stochastic Games (POSG) is challenging since each agent does not have access to the whole state information and, in case of concurrent learners, the environment has non-stationary dynamics. These problems could be partially overcome if the policies followed by the other agents were known, and, for this reason, many approaches try to estimate them through the so-called opponent modeling techniques. Although many researches have been devoted to the study of the accuracy of the estimation of opponents' policies, still little attention has been deserved to understand in which situations these model estimations can be actually useful to improve the agent's performance. This paper presents a preliminary study about the impact of using opponent modeling techniques to learn the solution of a POSG. Our main purpose is to provide a measure of the gain in performance that can be obtained by exploiting information about the policy of other agents, and how this gain is affected by the accuracy of the estimated models. Our analysis focus on a small two-agent POSG: the Kuhn Poker, a simplified version of classical poker. Three cases will be considered according to the agent knowledge about the opponent's policy: no knowledge, perfect knowledge, and imperfect knowledge. The aim is to identify which is the maximum error that can affect the model estimate without leading to a performance lower than that reachable without using opponent-modeling information. Finally, we will show how the results of this analysis can be used to improve the performance of a reinforcement-learning algorithm coped with a simple opponent modeling technique.

215B

Graph Laplacian Based Transfer Learning in Reinforcement Learning

Yi-Ting Tsao, Ke-Ting Xiao, Von Wun Soo

The aim of transfer learning is to accelerate learning in related domains. In reinforcement learning, many different features such as a value function and a policy can be transferred from a source domain to a related target domain. Many researches focused on transfer using hand-coded translation functions that are designed by the experts a priori. However, it is not only very costly but also problem dependent. We propose to apply the Graph Laplacian that is based on the spectral graph theory to decompose the value functions of both a source domain and a target domain into a sum of the basis functions respectively. The transfer learning can be carried out by transferring weights on the basis functions of a source domain to a target domain. We investigate two types of domain transfer, scaling and topological. The results demonstrated that the transferred policy is a better prior policy to reduce the learning time.

216C

Autonomous Agent Learning using an Actor-Critic Algorithm and Behavior Models

Victor Uc Cetina

We introduce a Supervised Reinforcement Learning (SRL) architecture for autonomous learning problems where an agent is required to deal with high dimensional spaces. Based on such architecture, we present an actor-critic algorithm. In our learning algorithm, behavior models learned from a set of examples, are used to dynamically reduce the set of relevant actions at each state of the environment encountered by the agent. Such subsets of actions are used to guide the agent through promising parts of the action space, avoiding the selection of useless actions. The algorithm handles continuous states and actions. Our experimental work with a difficult robot learning task shows clearly how this approach can significantly speed up the learning process and improve the final performance.

217A

Teaching Sequential Tasks with Repetition through Voice and Vision

Harini Veeraraghavan, Manuela Veloso

With the increasing number of applications of autonomous robots in human environments, robots are

beginning to interact and even collaborate with humans for solving various tasks. In this work, we present a multi-modal teaching by demonstration approach for teaching complex sequential tasks containing repetitions using voice and vision. Concretely, the human teaches by actively directing a humanoid robot to different locations using sound and providing visually interpretable cues. The robot in turn, uses these cues to actually perform the sequence of actions in order to accomplish the task during demonstration. The demonstrated task consists of a sequence of actions with repetitions resulting either from repeating sub-tasks or repetitions resulting from failed outcomes of the actions performed by the robot. By using the example demonstration, namely, the set of performed actions and the corresponding states, the robot learns a plan for executing the same task at a later time.

218B
Adaptive Kanerva-based Function Approximation for Multi-Agent Systems
Cheng Wu, Waleed Meleis
In this paper, we show how prototype optimization can be used to improve the performance of function approximation based on Kanerva Coding when solving large-scale instances of classic multi-agent problems. We apply our techniques to the predator-prey pursuit problem. We first demonstrate that Kanerva Coding applied within a reinforcement learner does not give good results. We then describe our new prototype optimization algorithm, based on prototype deletion and generation. We show that probabilistic prototype deletion with random prototype generation increases the fraction of test instances that are solved from 45% to *90%*, and that state splitting increases that fraction to *94%.* We also show that optimizing prototypes reduces the number of prototypes, and therefore the number of features, needed to achieve a *90%* solution rate by up to *87%*. These results demonstrate that our approach can dramatically improve the quality of the results obtained and reduce the number of prototypes required. We conclude that prototype optimization can greatly improve a Kanerva-based reinforcement learner's ability to solve large-scale multi-agent problems.

219C
Efficient Multi-Agent Reinforcement Learning through Automated Supervision
Chongjie Zhang, Sherief Abdallah, Victor Lesser
Multi-Agent Reinforcement Learning (MARL) algorithms suffer slow convergence and even divergence, especially in large-scale systems. In this work, we developed a supervision framework to speed up the convergence of MARL algorithms in a network of agents. The framework defines an organizational structure with automated supervision and a communication mechanism for exchanging information between lower-level agents and higher-level supervising agents. The abstracted states of lower-level agents travel upwards so that higher-level supervising agents generate a broader view of the state of the network. This broader view results in creating supervisory information which is passed down the hierarchy. We propose a generic extension to MARL algorithms that integrates supervisory information into the learning process, guiding agents' exploration of their state-action space. The simulation results in a distributed task allocation problem show that our proposed framework increases both the likelihood of convergence and the speed of convergence.

Agent Reasoning

158B
Continual Collaborative Planning for Mixed-Initiative Action and Interaction
Michael Brenner
Multiagent environments are often highly dynamic and only partially observable which makes deliberative action planning computationally hard. In many such environments, however, agents can take a more proactive approach and suspend planning for partial plan execution, especially for active information gathering and interaction with others. This paper presents a new algorithm for Continual Collaborative Planning (CCP) that enables agents to deliberately interleave planning, acting, perception and communication. Our implementation of CCP has been evaluated with MAPSIM, a tool that automatically generates multiagent simulations from formal multiagent planning (MAP) domains. For different such simulations, we show how CCP leads to collaborative planning and acting and, despite minimal linguistic capabilities, to fairly natural dialogues between agents.

159C
Supervision and Diagnosis of Joint Actions in Multi-Agent Plans
Roberto Micalizio, Pietro Torasso
The paper considers the problem of supervising the execution of a multi-agent plan (MAP) where actions are executed concurrently by a team of cooperating agents in a partially observable environment. The paper formalizes a distributed approach to the supervision task where each agent is responsible for supervising the actions it executes. A MAP is therefore decomposed into sub-plans, each of which is assigned to an agent of the team. Every agent supervises its own sub-plan by performing two main activities: the “monitoring”, which keeps track of the agent status at each time instant and detects anomalies (e.g. action failures) caused by the occurrence of unexpected events such as faults; and the “diagnosis”, which explains the detected anomalies in terms of agent failures.
In most MAPs there are causal dependencies among actions performed by different agents, since the agents have to cooperate by exchanging services or by executing joint-actions i.e., complex actions which require a number of agents to be completed. Therefore, the decomposition of the multi-agent plan in sub-plans does not produce, in general, sub-plans which are independent of each other; it follows that a fault affecting an agent can have harmful effects on the execution of plans assigned to other agents.
The paper describes how the monitoring of plan execution can be performed in a distributed way, even in presence of joint actions, by exploiting the notion of “dependency set”. A dependency set singles out, at each step during the plan execution, the minimal number of agents that have to cooperate to consistently supervise the plan execution.
Moreover, the paper introduces and formally characterizes the notions of agent and plan diagnosis; while plan diagnosis is able to capture the distinction between “primary failures” and “threatened actions”, the agent diagnosis provides information on the actual health status of the agents. The paper describes a strategy of failure propagation, which captures the interplay between agent diagnosis and plan diagnosis, and plays a critical

role in the understanding at what extent a fault affecting the functionalities of an agent affects the global MAP too.

160A
Theoretical and Experimental Results on the Goal-Plan Tree Problem
Patricia Shaw, Berndt Farwer, Rafael Bordini
Agents programmed in BDI-inspired languages have goals to achieve and a library of plans that can be used to achieve them, typically requiring further goals to be adopted. This is most naturally represented by a structure that has been called a Goal-Plan Tree. One of the uses of such structure is in agent deliberation (in particular, deciding whether to commit to achieving a certain goal or not). This paper presents new experimental results combining various types of goal-plan tree reasoning from the literature.

161B
Robust and Efficient Plan Recognition for Dynamic Multi-agent Teams
Gita Sukthankar, Katia Sycara
This paper addresses the problem of plan recognition for multi-agent teams. Complex multi-agent tasks typically require dynamic teams where the team membership changes over time. Teams split into subteams to work in parallel, merge with other teams to tackle more demanding tasks, and disband when plans are completed. We introduce a new multi-agent plan representation that explicitly encodes dynamic team membership and demonstrate the suitability of this formalism for plan recognition. From our multi-agent plan representation, we extract local temporal dependencies that dramatically prune the hypothesis set of potentially-valid team plans. The reduced plan library can be efficiently processed to obtain the team state history. Naive pruning can be inadvisable when low-level observations are unreliable due to sensor noise and classification errors. In such conditions, we eschew pruning in favor of prioritization and show how our scheme can be extended to rank-order the hypotheses. Experiments show that this robust pre-processing approach ranks the correct plan within the top 10%, even under conditions of severe noise.

Agent Cooperation

178A
Social Reward Shaping in the Prisoner's Dilemma
Monica Babes, Enrique Munoz, Michael Littman
Reward shaping is a common tool used to help reinforcement-learning agents converge more quickly to near-optimal behavior. In this paper, we introduce “social reward shaping”, which is reward shaping applied in the multiagent learning setting. We present preliminary experiments in the iterated Prisoner's dilemma setting that show that agents that use social reward shaping appropriately can behave more effectively than other classical learning and non-learning strategies. In particular, we show that these agents can both lead – encourage adaptive opponents to stably cooperate – and follow – adopt a best-response strategy to a fixed opponent – where classical approaches must choose between these behaviors.

179B
Mitigating Catastrophic Failure at Intersections of Autonomous Vehicles
Kurt Dresner, Peter Stone
Fully autonomous vehicles promise enormous gains in safety, efficiency, and economy for transportation. However, before such gains can be realized, a plethora of safety and reliability concerns must be addressed. Dresner and Stone have introduced a system for managing autonomous vehicles at intersections that is capable of handling more vehicles and causing fewer delays than modern-day mechanisms such as traffic lights and stop signs. While the authors claim that the system is safe, they do not discuss the possibility or implications of unforeseen mechanical failures. Because the system orchestrates such precarious “close calls” the tolerance for error is significantly diminished. In this paper, we make four main contributions. First, we perform a basic failure mode analysis on Dresner and Stone's mechanism, demonstrating that without any changes, it is unsuitable for use in the real world, due to its propensity for catastrophic failures. Second, we propose a method to mitigate these failures, bringing them into line with conventional intersection control mechanisms. Third, we give extensive empirical evidence suggesting that not only is this method effective, but that it is robust in the face of poor communications. Finally, we provide an analysis of the data indicating that despite the apparent potential for disastrous accidents, autonomous intersection management is likely to improve driver safety considerably.

180C
Multi-Agent Search using Sensors with Heterogeneous Capabilities
Guruprasad KR, Debasish Ghose
In this paper we introduce a new concept namely, *generalized Voronoi partition* and use it to formulate two heterogeneous multi-agent search strategies. The core idea is optimal deployment of agents having sensors with heterogeneous capabilities, in a search space so as to maximize search effectiveness. We address a few theoretical issues such as optimality of deployment, convergence and spatial distributedness of the control law and the search strategies.

181A
A new approach to cooperative pathfinding
Renee Jansen, Nathan Sturtevant
In the multi-agent pathfinding problem, groups of agents need to plan paths between their respective start and goal locations in a given environment, usually a two-dimensional map. Existing approaches to this problem include using static or dynamic information to help coordination. However, the resulting behaviour is not always desirable, in that too much information is hand-coded into the problem, agents take paths which look unintelligent, or because the agents collide and must re-plan frequently. We present a distributed approach in which agents share information about the direction in which they traveled when they passed through each location. This information is then used to encourage agents passing through the same location to travel in the same direction as previous agents. In addition to this new approach, we present performance metrics for multi-agent path planning as well as experimental results for the new approach. These results indicate that the number of collisions between agents is reduced and that the visual fidelity is improved.

182B
Resource constrained distributed constraint optimization using resource constraint free pseudo-

tree

Toshihiro Matsui, Marius Silaghi, Katsutoshi Hirayama, Makoto Yokoo, Hiroshi Matsuo

The Distributed Constraint Optimization Problem (DCOP) is a fundamental formalism for multi-agent cooperation. With DCOPs, the agent states and the relationships between agents are formalized into a constraint optimization problem, which is then solved using distributed cooperative optimization algorithms. In the original DCOP framework, a set of objective functions is employed to represent the relationships between agents. However, constraints for resources that are consumed by teams of agents are not well supported. Resource constraints are necessary to handle practical problems including distributed task scheduling with limited resource availability. A dedicated framework called Resource Constrained DCOP (RCDCOP) has been recently proposed. RCDCOP models objective functions and resource constraints separately. A resource constraint is an n-ary constraint which represents the limit on the number of resources of a given type available to agents. Previous research addressing RCDCOPs employs the Adopt algorithm, which is an efficient solver for DCOPs. An important graph structure for Adopt is the pseudo-tree for constraint networks. A pseudo-tree implies a partial ordering of variables. In this variable ordering, n-ary constrained variables are placed on a single path of the tree. Therefore, resource constraints that have large arity augment the depth of the pseudo-tree. This also reduces the parallelism, and therefore the efficiency of Adopt. In this paper we propose another version of the Adopt algorithm for RCDCOP using a pseudo-tree which is generated ignoring resource constraints. The key ideas of our work are as follows: (i) The pseudo-tree is generated ignoring resource constraints. (ii) Virtual variables are introduced, representing the usage of resources. These virtual variables are used to share resources among sub-trees. (iii) The addition of virtual variables increases the search spaces. To reduce this problem, the search is pruned using the bounds defined by the resource constraints. These ideas are used to extend Adopt. The proposed method reduces the previous limitations in the construction of RCDCOP pseudo-trees. The efficiency of our technique depends on the class of problems being considered, and we describe the obtained experimental results.

183C

Coordination of First Responders Under Communication and Resource Constraints

Robert Lass, Joseph Kopena, Evan Sultanik, Duc Nguyen, Christopher Dugan, William Regli

Coordination between first responders is a key requirement in successfully managing a natural or other disaster. At all levels of the task, from integrating heterogeneous systems to addressing response tasks and allocating resources, decisions must be made in globally optimal fashions or risk inefficient and potentially counter-productive efforts. Automated coordination mechanisms stand to benefit emergency personnel in sharing information and arriving at global solutions. However, challenges of network communication in these settings must be addressed to make such mechanisms effective. This paper discusses the application of distributed constraint optimization to coordination in disaster management. It presents an implemented example system for the problem of shelter assignment and outlines some of the networking challenges and future research directions that must be addressed before real-world use of distributed constraint optimization becomes a reality.

184A

Replacing the Stop Sign: Unmanaged Intersection Control for Autonomous Vehicles

Mark Van Middlesworth, Kurt Dresner, Peter Stone

As computers inevitably begin to replace humans as the drivers of automobiles, our current human-centric traffic management mechanisms will give way to hyper-efficient systems and protocols specifically designed to exploit the capabilities of fully autonomous vehicles. Dresner and Stone have introduced such a system for coordinating large numbers of autonomous vehicles at intersections. Their experiments suggest that this system could alleviate many of the dangers and delays associated with intersections by allowing vehicles to ``call ahead'' to an agent stationed at the intersection and reserve time and space for their traversal. Unfortunately, such a system is not cost-effective at small intersections, as it requires the installation of specialized infrastructure. In this paper, we propose an intersection control mechanism for autonomous vehicles designed specifically for low-traffic intersections where a system such as Dresner and Stone's would not be practical, just as inexpensive stop signs are used at intersections that do not warrant a full traffic light installation. Our mechanism is based on purely peer-to-peer communication and thus requires no infrastructure at the intersection. We present experimental results demonstrating that our system, while not suited to large, busy intersections, can significantly outperform traditional stop signs at small intersections: vehicles spend less time waiting and consume less fuel.

185B

An Improved Dynamic Programming Algorithm for Coalition Structure Generation

Talal Rahwan, Nick Jennings

Forming effective coalitions is a major research challenge in the field of multi-agent systems. Central to this endeavour is the problem of partitioning the set of agents into exhaustive and disjoint coalitions such that the social welfare is maximized. This coalition structure generation problem is extremely challenging due to the exponential number of partitions that need to be examined. Specifically, given n agents, there are $O(n^n)$ possible partitions. To date, the only algorithm that can find an optimal solution in $O(3^n)$ is the Dynamic Programming (DP) algorithm, due to Rothkopf et al. However, one of the main limitations of DP is that it requires a significant amount of memory. In this paper, we devise an Improved Dynamic Programming algorithm (IDP) that is proved to perform fewer operations than DP (e.g. 38.7% of the operations given 25 agents), and is shown to use only 33.3% of the memory in the best case, and 66.6% in the worst.

186C

Discovering Tactical Behaviour Patterns Supported by Topological Structures in Soccer-Agent Domains

Fernando Ramos, Huberto Ayanegui

This work deals with the discovery of tactical plays in soccer-agent domains. Due to the nature of team work in soccer-agent domains, the discovery of tactical plays should take into account the players involved in such plays within the context of team formations. Nevertheless, the dynamic nature and the multiple interactions between players at each instant of the game difficult the tracking of formations, which in turn difficult

the discovery of tactical plays. The work presented in this paper is supported by an efficient tracking of formations, even though they are submitted to dynamic changes of the world, based on the construction of topological structures. A generalization mechanism is applied to a set of tactical plays extracted from real matches played by teams of the Robocup simulation league. Successful results show that the proposed model is able to discover tactical behaviour patterns which are mainly characterized by the path of the ball, the agents participating in the play and the zones of the field where the plays have taken place.

187A

A Best-First Anytime Search Algorithm for Coalition Structure Generation

Chattrakul Sombattheera, Aditya Ghose

This work presents a best-first anytime search algorithm for optimal coalition structures. The algorithm is based on a novel approach that it generates coalition structures based on coalition values while existing algorithms generate based on the structure (members and configurations) of coalitions. With our algorithm, coalition structures are generated by repeatedly choosing the best coalition, whose agent contribution to coalition structure is the highest, from available candidates. We carried out experiments to explore the performance of our algorithm in 20 data distributions. Empirical results show that our algorithm always converges the optimal coalition structures very quickly compared to the fastest algorithm. Although our algorithm terminates later (because of a simple prune mechanism being used) in half the cases, our algorithm always yields a better, or, at least, as good as the fastest algorithm at any point in time.

188B

RIAACT: A robust approach to adjustable autonomy for human-multiagent teams

Nathan Schurr, Janusz Marecki, Milind Tambe, Chien-Ju Ho, Jane Yung-jen Hsu

When human-multiagent teams act in real-time uncertain domains, adjustable autonomy (dynamic transferring of decisions between human and agents) raises three key challenges. First, the human and agents may differ significantly in their worldviews, leading to inconsistencies in their decisions. Second, these human-multiagent teams must operate and plan in real-time with deadlines with uncertain duration of human actions. Thirdly, adjustable autonomy in teams is an inherently distributed and complex problem that cannot be solved optimally and completely online. To address these challenges, our paper presents a solution for Resolving Inconsistencies in Adjustable Autonomy in Continuous Time (RIAACT). RIAACT incorporates models of the resolution of inconsistencies, continuous time planning techniques, and hybrid method to address coordination complexity. These contributions have been realized in a disaster response simulation system.

189C

Adaptive Manager-side Control Policy in Contract Net Protocol for Massively Multi-Agent Systems

Sugawara Toshiharu, Satoshi Kurihara, Toshio Hirotsu, Kensuke Fukuda

We describe a new adaptive manager-side control policy for the contract net protocol for a massively multi-agent system (MMAS). To improve overall performance of MMAS, tasks must be allocated to appropriate agents. From this viewpoint, a number of negotiation protocols were proposed in the MAS context, but most assume a small-scale, unbusy environment. We previously reported that, using contract net protocol (CNP), the overall efficiency could improve by an adequate control of degree of fluctuation in the awarding phase depending on the state of MMAS. In this paper, we propose the method to estimate these states from the bid values, which have hitherto not been used effectively. Then the manager-side policy flexibly and autonomously with some degree of fluctuation responsive to the estimated states is introduced. We also evaluate that our proposed CNP policy.

190A

Towards Bidirectional Distributed Matchmaking

Victor Shafran, Gal Kaminka, Sarit Kraus, Claudia Goldman

Matchmaking is the process of introducing two or more agents to each other. Current matchmaking techniques are unidirectional and fail to address large-scale and highly dynamic systems with time constraints. We propose a new distributed technique which scales well, and still maintains relatively low matchmaking time and communication overhead. Our technique introduces very low storage and computational overhead to the agents. We suggest using a matching cache which can take advantage of the multidirectional nature of the matchmaking problem. We empirically evaluate the proposed technique on bilateral matchmaking and show that it outperforms the existing techniques.

191B

Designing Human-Computer Multi-agent Collaboration in Productive Multi-Player Games

Wenn-Chieh Tsai, Yuan-Hsiang Lee, Tsung-Hsiang Chang

This research explores productive multi-player games as a platform for human-computer agent collaboration. A multiagent perspective is taken to examine the principles of both gameplay and mechanism design for productive games. To engage human players in sustained gameplay, the game agents are designed with the flow and dramatic principles. To ensure productivity, the game mechanism is designed such that rational agents, both human and software, will follow the target strategy to reach subgame perfect equilibrium. The design principles are demonstrated and evaluated using PhotoSlap, a multi-player productive game for photo annotation.

192C

Trading Off Solution Quality for Faster Computation in DCOP Search Algorithms

William Yeoh, Sven Koenig, Xiaoxun Sun

Distributed Constraint Optimization Problems (DCOPs) have been known to be useful in modeling multiagent coordination problems. However, its drawback is its inefficiency in finding optimal solutions, especially for large problems. For these computationally challenging problems, we propose two mechanisms that allow tradeoffs between solution quality and computation time for two DCOP search algorithms, ADOPT and BnB-ADOPT. The primary contribution of this paper is that the final solution quality found using the proposed mechanisms will always be bounded by a more interpretable and meaningful error bound when compared with the existing mechanism. Additionally, it is shown empirically that one of the mechanisms dominates the existing mechanism in BnB-ADOPT and is no worse than the existing mechanism in ADOPT.

193A
Anytime Local Search for Distributed Constraint Optimization
Roie Zivan
Most former studies of Distributed Constraint Optimization Problems (DisCOPs) search considered only complete search algorithms, which are practical only for relatively small problems. Distributed local search algorithms can be used for solving DisCOPs. However, because of the differences between the global evaluation of a system's state and the private evaluation of states by agents, agents are unaware of the global best state which is explored by the algorithm. Previous attempts to use local search algorithms for solving DisCOPs reported the state held by the system at the termination of the algorithm, which was not necessarily the best state explored. A general framework for implementing distributed local search algorithms for DisCOPs is proposed. The proposed framework makes use of a BFS-tree in order to accumulate the costs of the system's state in its different steps and to propagate the detection of a new best step when it is found. The resulting framework enhances local search algorithms for DisCOPs with the anytime property. The proposed framework does not require additional network load. Agents are required to hold a small (linear) additional space (beside the requirements of the algorithm in use). The proposed framework preserves privacy at a higher level than complete DisCOP algorithms which make use of a pseudo-tree (ADOPT, DPOP).

Agent Theories, Models and Architectures

106A
Reasoning about agent deliberation
Natasha Alechina, Mehdi Dastani, Brian Logan, John-Jules Meyer
We present a family of sound and complete logics for reasoning about execution strategies for SimpleAPL programs. SimpleAPL is a fragment of the agent programming language 3APL designed for the implementation of cognitive agents with beliefs, goals and plans. The logics are variants of PDL, and allow us to prove safety and liveness properties of agent programs for both normal executions and `abnormal' executions where the preconditions of a basic action fail. We show how to axiomatize different program execution strategies for SimpleAPL programs, and, for each execution strategy, prove a correspondence between the operational semantics of SimpleAPL and the models of the corresponding logic. We illustrate the utility of our approach with an example in which we show how to verify correctness properties for a simple agent program under different execution strategies.

107B
Negotiation by Induction
Chiaki Sakama
This paper presents a logical framework for automated negotiation. An agent accepts a proposal if it is proved by its knowledge base. If this is not the case, an agent seeks conditions to accept a proposal or may give up some of its current belief to reach an agreement. These attitudes of agents are characterized using induction and default reasoning.

108C
Epistemic Logic and Explicit Knowledge in Distributed Programming
Andreas Witzel, Jonathan Zvesper
In this paper we propose an explicit form of knowledge-based programming. Our initial motivation is the distributed implementation of game-theoretical algorithms, but we abstract away from the game-theoretical details and describe a general scenario, where a group of agents each have some initially private bits of information which they can then communicate to each other. We draw on existing literature to give a formal model using modal logic to represent the knowledge of the agents as well as how that knowledge changes as they communicate. We sketch an implementation which enables processes in a distributed system to explicitly evaluate knowledge formulae. Then we prove that the implementation captures the formal model, and therefore correctly reflects the general scenario. Finally we look at how our approach lends itself to generalisations, and discuss application perspectives.

Agent Societies and Societal Issues

120C
Identifying Beneficial Teammates using Multi-Dimensional Trust
Jaesuk Ahn, Xin Sui, David DeAngelis, Suzanne Barber
Multi-agent teams must be capable of selecting the most beneficial teammates for different situations. Multi-dimensional trustworthiness assessments have been shown significantly beneficial to agents when selecting appropriate teammates to achieve a given goal. Reliability, quality, availability, timeliness and compatibility define the behavioral constraints of the multi-dimensional trust (MDT) model. Given the MDT model in this research, an agent learns to identify the most beneficial teammates by prioritizing each dimension differently. An agent's attitudes towards rewards, risks and urgency are used to drive an agent's prioritization of dimensions in a MDT model. Each agent is equipped with a Temporal-Difference (TD) learning mechanism with tile coding to identify its optimal set of attitudes and change its attitudes when the environment changes. Experimental results show that changing attitudes to give preferences for respective dimensions in the MDT offers a superior means to finding the best teammates for goal achievement.

121A
Simulating the Effects of Sanction for the Emergence of Cooperation in a Public Goods Game
Ana Bazzan, Silvio Dahmen, Alexandre Baraviera
Several explanations have been proposed in order to explain why, in public goods games, cooperation does not collapse. In these games free-riders enjoy the benefits of other individuals who contribute in benefit of a community. In the present work we address a public goods game where individuals have the choice between contributing to a sanctioning institution and to a sanction-free one. In the former there is a possible sanction for those who do not contribute. Our results show that individuals who contribute to a sanctioning institution are better off after several repetitions of the game, despite the costs associated with sanctioning. This reproduces results found in experiments with human subjects, which point to advantages of sanctioning measures as a factor for the stabilization of cooperation.

123C
Learning Task-Specific Trust Decisions
Ikpeme Erete, Erin Ferguson, Sandip Sen
We study the problem of agents locating other agents that are both capable and willing to help complete assigned tasks. An agent incurs a fixed cost for each help request it sends out. To minimize this cost, the performance metric used in our work, an agent should learn based on past interactions to identify agents likely to help on a given task. We compare three trust mechanisms: success-based, learning-based, and random. We also consider different agent social attitudes: selfish, reciprocative, and helpful. We evaluate the performance of these social attitudes with both homogeneous and mixed societies. Our results show that learning-based trust decisions consistently performed better than other schemes. We also observed that the success rate is significantly better for reciprocative agents over selfish agents.

124A
Specifying and Enforcing Norms in Artificial Institutions
Nicoletta Fornara, Marco Colombetti
In this paper we investigate two related aspects of the formalization of open interaction systems: how to specify norms, and how to enforce them by means of sanctions. The problem of specifying the sanctions associated with the violation of norms is crucial in an open system because, given that the compliance of autonomous agents to obligations and prohibitions cannot be taken for granted, norm enforcement is necessary to constrain the possible evolutions of the system, thus obtaining a degree of predictability that makes it rational for agents to interact with the system. In our model, we introduce a construct for the definition of norms in the design of artificial institutions, expressed in terms of roles and event times, which, when certain activating events take place, is transformed into commitments of the agents playing certain roles. Norms also specify different types of sanctions associated with their violation. In the paper, we analyze the concept of sanction in detail and propose a mechanism through which sanctions can be applied.

126C
An Adaptive Probabilistic Trust Model and its Evaluation
Chung-Wei Hang, Yonghong Wang, Munindar Singh
In open settings, the participants are autonomous and there is no central authority to ensure the felicity of their interactions. When agents interact in such settings, each relies upon being able to model the trustworthiness of the agents with whom it interacts. Fundamentally, such models must consider the past behavior of the other parties in order to predict their future behavior. Further, it is sensible for the agents to share information via referrals to trustworthy agents. Much progress has recently been made on probabilistic trust models including those that support the aggregation of information from multiple sources. However, current models do not support trust updates, leaving updates to be handled in an ad hoc manner.
This paper proposes a trust representation that combines probabilities and certainty (defined as a function of a probability-certainty density function). Further, it offers a trust update mechanism to estimate the trustworthiness of referrers. This paper describes a testbed that goes beyond existing testbeds to enable the evaluation of a composite probability-certainty model. It then evaluates the proposed trust model showing that the trust model can (a) estimate trustworthiness of damping and capricious agents correctly, (b) update trust values of referrers accurately, and (c) resolve the conflicts in referral networks by certainty discounting.

127A
Extending Virtual Organizations to improve trust mechanisms
Ramón Hermoso, Roberto Centeno Sánchez, Holger Billhardt, Sascha Ossowski
Virtual Organizations (VOs) are becoming more and more important as a field of researching in Multi-Agent Systems (MAS). The problem of selecting suitable counterparts to interact with is of particular relevance for agents belonging to a VO. This issue has been extensively investigated, applying probability or cognitive approaches but very few focus has been given to the use of internal organizational structures and the improvement they can provide. In this paper we analyze how organizational structures can support the agent selection process based an trust mechanisms. Furthermore, we present a way to extend VOs automatically (e.g., their role taxonomies) by detecting and identifying new roles. We show that such extensions lead to an improvement of agents decisions when employing trust mechanisms that take advantage of organizational structures.

129C
Convergence at Prominent Agents: A Non-Flat Synchronization Model of Situated Multi-Agents
Jiuchuan Jiang, Yichuan Jiang
This paper presents a novel non-flat synchronization model where the synchronization capacity of each agent is different regarding its social rank and strategy dominance. In the presented model, the prominent agents may have higher synchronization forces, and finally the collective synchronization results may incline to converge at such prominent agents' strategies, which is called prominence convergence in collective synchronization and proved by our experimental results. The presented model can well match the peculiarities of real multi-agent societies where each agent plays a different role in the synchronization, and make up the restrictions of related benchmark works that only concerned about the flat synchronization.

130A
How Automated Agents Treat Humans and Other Automated Agents in Situations of Inequity: An Experimental Study
Ron Katz, Sarit Kraus
This paper explores the question of how agent designers perceive and treat their agent's opponents. In particular, it examines the influence of the opponent's identity (human vs. automated agent) in negotiations. We empirically demonstrate that when people interact spontaneously they treat human opponents differently than automated agents in the context of equity and fairness considerations. However, this difference vanish when people design and implement agents that will interact on their behalf. Nevertheless, the commitment of the agents to honor agreements with people is higher than their commitment to other agents. In the experiments, which consisted of 147 computer science students, we used the Colored Trails game as the negotiation environment. We suggest possible explanations for the relationships among online

players, agent designers, human opponents and automated opponents.

132C
A Distributed Normative Infrastructure for Situated Multi-Agent Organisations
Fabio Okuyama, Rafael Bordini, Antonio Costa
In most of the existing approaches to the design of multi-agent systems, there is no clear way in which to relate organisational and normative structures to the model of the environment where they are to be situated and operate. Our work addresses this problem by putting together, in a practical approach to developing multi-agent systems (and social simulations in particular), a high-level environment modelling language that incorporates aspects of agents, organisations, and normative structures. The paper explains in some detail how the ideas of normative objects and normative places, put together as a distributed normative infrastructure, allow the definition of certain kinds of situated multi-agent organisations, in particular organisations for multi-agent systems that operate within concrete environments. Normative objects are environment objects used to explicitly convey normative content that regulate the behaviour of agents within the place where such objects can be perceived by agents. The paper briefly introduces such concepts, showing how they were integrated into the MAS-SOC multi-agent systems platform for social simulation, and hints on new problems of (situated) organisational and normative structures that were brought forward by the work presented here.

133A
Modeling and Managing Collective Cognitive Convergence
Van Parunak, Ted Belding, Rainer Hilscher, Sven Brueckner
When the same set of people interact frequently with one another, they grow to think more and more along the same lines, a phenomenon we call "collective cognitive convergence" (C3). In this paper, we discuss instances of this phenomenon and why it is advantageous or disadvantageous; review previous work in sociology, computational social science, and evolutionary biology that sheds light on C3; define a computational model for the convergence process and quantitative metrics that can be used to study it; report on experiments with this model and metrics; and suggest how the insights from this model can inspire techniques for managing C3.

135C
Determining Top K Nodes in Social Networks using the Shapley Value
Narayanam RamaSuri, Y. Narahari
In this paper, we consider the problem of selecting, for any given positive integer k, the top-k nodes in a social network, based on a certain measure appropriate for the social network. This problem is relevant in many settings such as analysis of co-authorship networks, diffusion of information, viral marketing, etc. However, in most situations, this problem turns out to be NP-hard. The existing approaches for solving this problem are based on approximation algorithms and assume that the objective function is submodular. In this paper, we propose a novel and intuitive algorithm based on the Shapley value, for efficiently computing an approximate solution to this problem. Our proposed algorithm does not use the sub-modularity of the underlying objective function and hence it is a general approach. We demonstrate the efficacy of the algorithm using a co-authorship data set from e-print arXiv (www.arxiv.org), having 8361 authors.

136A
Simulating Human Behaviors in Agent Societies
Alicia Ruvinsky, Michael Huhns
As increasing numbers of processors and agents pervade the human environment, societies comprising both humans and agents will emerge. Presently, it is unknown how a person might fair in such mixed societies. For the societies to operate effectively and efficiently, it is important for the humans and agents to recognize and understand each other's behavior. This paper provides an initial step in that understanding, via two contributions: (1) we provide models, within a limited domain, for agents that behave like humans and (2) we present the results of simulated interactions between the human-like agents and a variety of purely rational agents. Our models for the behaviors of people are based on recent sociological research by Simpson and Willer that explores the motivation for humans' cooperative pro-social behavior, a conceivably non-rational process. Modeling human behaviors presents a means of exploring and understanding motivations, consequences, and resolutions to human-agent interactions. We aspire to exploit this knowledge about human behavior in order to observe its ramifications in an agent world, and to motivate development of human-agent societies. Our results show that, although there are pitfalls to which humans are vulnerable, there exist niches for human prosperity in a rational agent world.

137B
Do Humans Identify Efficient Strategies in Structured Peer-to-Peer Systems? (Short Paper)
Stephan Schosser, Klemens Böhm, Bodo Vogt
In the last years, distributed coordinator-free systems, e.g., peer-to-peer systems (P2P systems), have attracted much interest among researchers and practitioners. In these systems it is difficult to motivate participants to cooperate. To this end, researchers have proposed various incentive mechanisms. In this paper we are interested in the following question: Do human beings indeed use the strategies that are rational in presence of the incentive mecha-nism? As humans control the agents in distributed coordinator-free systems, e.g., the peers in peer-to-peer systems, answering this question is essential. We conduct human experiments in the context of structured P2P systems to answer it. This paper shows that humans tend to find it difficult to resort to the strategies ex-pected by the system designer.

138C
Social norm emergence in virtual agent societies
Bastin Tony Roy Savarimuthu, Maryam Purvis, Martin Purvis
The advent of virtual environments such as SecondLife call for a distributed approach for norm emergence and spreading. In open virtual environments, monitoring various interacting agents (avatars), using a centralized authority might be computationally expensive. The number of possible states and actions of an agent could be huge. An approach for sustaining order and smoother functioning of these environments can be facilitated through norms. Agents can generate norms based on interactions. In particular, those social norms that incur certain cost to an individual agent but benefit the whole society are more interesting than those benefit both the

agent and the society. The problem is that the selfish agents might not be willing to share the norm adherence cost. In this work, we experiment with notion proposed by Axelrod that social norms are best at preventing small defections where the cost of enforcement is low. We also study how common knowledge can be used to faciliate the overall benefit of the society. We believe our work can be used to facilitate norm emergence in virtual online societies.

139A
Searching and Sharing Information In Networks of Heterogeneous Agents
George Vouros
Finding the right agents in a large and dynamic network to provide the needed resources in a timely fashion, is a long standing problem. This paper extends previous work on a tuning method for information searching and sharing in large-scale and dynamic networks of agents and studies the effectiveness of the method when agents are heterogeneous: I.e. when agents do not share a common conceptualization and when agents have different preferences on specific information categories. The extended method enables agents to search effectively by building overlay networks for specific information categories maintaining shortcuts with non-neighbors, by acquiring their neighbors' interests, advertising their information provision abilities and maintaining indices for routing queries, in an integrated way. Specifically, the paper demonstrates through performance experiments how networks of heterogeneous agents can be "tuned" to answer queries by gathering information about their acquaintances and by imposing overlay structures to the network of acquaintances, even if their expertise shifts.

140B
Stable cooperation in changing environments
Humphrey Sorensen
This paper addresses the issue of emergence of robust cooperation among self-interested agents interacting in N-player social dilemma games. These games represent a useful abstraction of many types of agent interactions. Much research has focused on two player games while N-player extensions are, in comparison, relatively under-represented in the literature.
Past research has attempted to understand the effect of spatial constraints on a society or population of agents. Much of this work has focussed on grid-like topologies to constrain the interactions of the agents. Recently, there has been an increased interest in modelling more realistic spatial topologies including random graphs and small world networks.
In this paper, we enforce a regular weighted graph topology on the agents such that the graph exhibits a high degree of community structure. In this work, a series of graphs are created each exhibiting a different level of community structure; we show the influence that community structure has on the emergence of cooperation. In our simulations, agents interact with those agents with whom they are directly connected; the probability of interaction is dependant on the weight of the edge. Agents adopt new strategies by imitating neighbouring successful agents. Two simple adoption techniques are utilised.
Many open agent systems may have a high degree of uncertainty due to a number of factors including: agents may perform their acts incorrectly or imperfectly; their acts may be misinterpreted; agents may learn or imitate from others or agents may exit or join the group thereby changing the environment.
In this paper, we model two classes of agent strategy: i)a strategy set that involves a simple cooperation or defection and ii) a strategy set that represents a form of generalised tit-for-tat.
We also model two types of uncertainty: a simple form of noise where certain actions are mis-implemented and a more drastic change to the environment where in effect the payoffs in the game are reversed.
We show that for both sets of strategies, given a suitably high level of community structure in the spatial topology: i) cooperation will emerge ii) the cooperation is robust in the face of noise iii) the population can change behaviours in the face of dramatic environmental change, provided there exists some noise in the system.

Agent Communication

75C
Efficient Approximate Inference in Distributed Bayesian Networks for MAS-based Sensor Interpretation
Norman Carver
Distributed problem solving (DPS) is the subfield of multi-agent systems that is concerned with solving large-scale, often inherently distributed problems, using systems of distributed intelligent agents. One important application is sensor interpretation (SI) in distributed sensor networks. SI domains can frequently be modeled with Bayesian networks (BNs), with the interpretation process involving BN inference. Distributed, multi-agent SI can be modeled with distributed Bayesian networks (DBNs), where sub-networks of the global BN are distributed to different agents. In DBNs, some inferences require communication among the agents. The multiply sectioned Bayesian network (MSBN) framework is the most studied approach for DBN inference in an MAS setting. However, we do not believe the MSBN framework is well suited for large-scale MAS-based SI. In such problems, exact inference will be impractical, it will be critical to take advantage of the parallel computational capabilities of the agents, and agent coordination strategies may need to be flexible and dynamic. Among the key problems we see with the MSBN are: it supports only exact inference, its global propagation procedure reduces agents' ability to work in parallel, agent autonomy and asynchrony is limited, and communication patterns among the agents are severely restricted. This paper describes the elements of an alternative framework for inference in DBNs, which we have developed to support efficient, approximate MAS-based SI. Compared to the MSBN approach, our approach supports more autonomy and asynchrony among the agents, and more focused, situation-specific communication patterns. Its use can lead to significant improvements in agent utilization and time-to-solution performance. The paper presents some results from analyses of basic sample strategies to demonstrate that an MAS could potentially use our framework to produce acceptable quality SI solutions faster than with the MSBN.

76A
Synchronization Protocols for Reliable Communication in Fully Distributed Agent Systems
Hywel Dunn-Davies, Jim Cunningham, Shamima Paurobally

In order to prevent misunderstandings within groups of interacting agents, it is necessary to ensure that the agents' beliefs regarding the overall state of the interaction are consistent with each other at all times. Paurobally et al. (2003) proposed that these beliefs could be synchronized by adding a specialized protocol layer that incorporates protocols specifically designed to synchronize the agents' beliefs. Here we define the problem that such protocols would need to solve in the worst case, and prove it to be insoluble. We then consider the possibility of synchronizing the beliefs of groups of agents if it is assumed that the communication layer notifies the sender of a message whenever that message is not successfully delivered. Paurobally et al. (2003) proved that this assumption allows agents' beliefs to be synchronized in bilateral interactions. However, we prove that this assumption is insufficient to achieve belief synchronization in groups of three or more agents. Finally, we discuss the possibility of achieving adequate synchronization using probabilistic protocols.

77B
A Multi-Agent Based Implementation of a Delphi Process
Ivan Garcia-Magariño, Jorge Gomez Sanz, Jose R. Perez-Aguera
The Delphi protocol is applied when a community of experts is required to reach a consensus and to deliver an answer. In these cases, consensus stands for reaching an agreement among the experts about what the answer should be. This consensus reaching problem has been already considered in the literature, though its automatisation remains as a challenge. Intuitively, the experts should dialogue, interchange ideas, and change their mind as the discussion progresses. This paper presents the first complete-implementation of the Delphi process. This implementation is achieved with a Multi-agent System(MAS), in which the experts are implemented with agents. The presented case study solves the document relevance evaluation problem where a community of experts decide whether a document is relevant or not. In conclusion, this paper makes an important contribution to people using Delphi processes, because the presented system is the first complete-computerised Delphi process. With respect to multi-agent systems, it has the potential to solve coordination in an original way.

Economic Paradigms

254B
Strategyproof Deterministic Lotteries under Broadcast Communications
Alon Altman, Moshe Tennenholtz
The design of deterministic and fair mechanisms for selection among a set of self-motivated agents based solely on these agents' input is a major challenge in multiagent systems. This challenge is especially difficult when the agents can only communicate via a broadcast channel. We propose the notion of selection games: a special case of zero-sum games where the only possible outcomes are selections of a single agent among the set of agents. We assume the lack of an external coordinator, and therefore we focus on mechanisms which have a solution where the agents play weakly dominant strategies. Our first major result shows that dominated strategies could be added to any selection mechanism, so that the resulting mechanism becomes quasi-symmetric. For fairness, we require the mechanism to be non-imposing; that is, the mechanism should allow any agent to be selected in such a solution. We first show that such mechanisms do not exist when there are two or three agents in the system. However, surprisingly, we show that such mechanisms exist when there are four or more agents. Moreover, in our second major result, we show that there exist selection mechanisms that implement any distribution over the agents, when the agents play mixed dominant strategies. These results also have significance for electronic commerce, ranking systems, and social choice.

255C
Decommitment in Multi-resource Negotiation
Bo An, Victor Lesser, Kwang Mong Sim
This paper presents the design and implementation of negotiation agents that negotiate with other entities for acquiring multiple resources. Negotiation agents in this paper are designed to adjust 1) the number of tentative agreements for each negotiation and 2) the amount of concession they are willing to make in response to changing market conditions and negotiation situations. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each negotiation issue is dynamically determined by 1) the likelihood that negotiation will not be successfully completed (conflict probability), 2) the expected agreement price of the issue, and 3) the expected number of final agreements. The negotiation deadline of each negotiation is determined by its scarcity. Agents are permitted to decommit from agreements, and an agent can make more than one tentative agreement for each issue. The maximum number of tentative agreements made by an agent is constrained by the market situation and budget. Results from a series of experiments indicate that on average, our negotiation strategy achieved higher average utility than traditional negotiation strategies.

256A
Incentives in Effort Games
Yoram Bachrach, Jeffrey Rosenschein
We consider Effort Games, a game theoretic model of cooperation in open environments, which is a variant of the principal-agent problem from economic theory. In our multiagent domain, a common project depends on various tasks; achieving certain subsets of the tasks completes the project successfully, while others do not. The probability of achieving a task is higher when the agent in charge of it exerts effort, at a certain cost for that agent. A central authority, called the principal, attempts to incentivize agents to exert effort, but can only reward agents based on the success of the entire project.
We model this domain as a normal form game, where the payoffs for each strategy profile are defined based on the different probabilities of achieving each task and on the boolean function that defines which task subsets complete the project and which do not. We view this boolean function as a simple coalitional game, and call this game the underlying coalitional game. We show that finding the minimal reward that induces an agent to exert effort is at least as hard computationally as finding the Banzhaf power index in the underlying coalitional game, so this problem is #P-hard in general.
We also show that in a certain restricted domain, where the underlying coalitional game is a unanimity weighted voting game with certain properties, it is possible to solve all of the above problems in polynomial time.

257B

Learn While You Earn: Two Approaches to Learning Auction Parameters in Take-it-or-leave-it Auctions

Archie Chapman, Alex Rogers, Nick Jennings

Much of the research in auction theory assumes that the auctioneer knows the distribution of participants' valuations with complete certainty. However, this is unrealistic. Thus, we analyse cases in which the auctioneer is uncertain about the valuation distributions; specifically, we consider a repeated auction setting in which the auctioneer can learn these distributions. Using take-it-or-leave-it auctions (Sandholm and Gilpin, 2006) as an exemplar auction format, we consider two auction design criteria. Firstly, an auctioneer could maximise expected revenue each time the auction is held. Secondly, an auctioneer could maximise the information gained in earlier auctions (as measured by the Kullback-Liebler divergence between its posterior and prior) to develop good estimates of the unknowns, which are later exploited to improve the revenue earned in the long-run. Simulation results comparing the two criteria indicate that setting offers to maximise revenue does not significantly detract from learning performance, but optimising offers for information gain substantially reduces expected revenue while not producing significantly better parameter estimates.

258C

Incorporating User Utility Into Sponsored-Search Auctions

Yagil Engel, David Maxwell Chickering

We study principled methods for incorporating user utility into the selection of sponsored search advertisements; we argue that including the user in the allocation mechanism is critical to the long-term profitability of the search engine. We present a model of user experience using a structured utility function, and we present reasonable assumptions that simplify the assessment of this function. We describe simple controls, in the form of multiplicative and additive factors, which allow the publisher (i.e., the search engine) to trade off user utility with short-term revenue. We examine variations of GSP mechanisms that accommodate user-utility functions, and show that for those functions that are additive over advertisements, some of the theoretical properties in the traditional GSP mechanisms have interesting and useful parallels in our new variations. We provide simulation results that exemplify the difference between different choices for the utility model, and in particular we show that an additive tradeoff factor is a simple and cheap way to improve advertisement relevance.

259A

Nonuniform Bribery

Piotr Faliszewski

We study the concept of bribery in the situation where voters are willing to change their votes as we ask them, but where their prices depend on the nature of the change we request. Our model is an extension of the one of Faliszewski et al. [FHH06], where each voter has a single price for any change we may ask for. We show polynomial-time algorithms for our version of bribery for a broad range of voting protocols, including plurality, veto, approval, and utility based voting. In addition we prove NP-completeness for a couple of our nonuniform bribery problems for weighted voters, and give approximation algorithms for two NP-complete bribery problems defined in [FHH06].

260B

A Preliminary result on a representative-based multi-round protocol for multi-issue negotiations

Katsuhide Fujita, Takayuki Ito, Mark Klein

Multi-issue negotiation protocols represent a promising field since most negotiation problems in the real world involve multiple issues. Our work focuses on negotiation with interdependent issues, in which agent utility functions are nonlinear. Existing works have not yet focused on agents' private information. In addition, they were not scalable in the sense that they have shown a high failure rate for making agreements among 5 or more agents. In this paper, we focus on a novel multi-round representative-based protocol that utilizes the amount of agents' private information revealed. Experimental results demonstrate that our mechanism reduces the failure rate in making agreements, and it is scalable on the number of agents compared with existing approaches.

261C

Approximating Mixed Nash Equilibria using Smooth Fictitious Play in Simultaneous Auctions

Enrico Gerding, Andrew Byde, Edith Elkind, Zinovi Rabinovich, Nick Jennings

We investigate equilibrium strategies for bidding agents that participate in multiple, simultaneous second-price auctions with perfect substitutes. For this setting, previous research has shown that it is a best response for a bidder to participate in as many such auctions as there are available, provided that other bidders only participate in a single auction. In contrast, in this paper we consider equilibrium behaviour where all bidders participate in multiple auctions. For this new setting we first prove that there exist no pure Nash equilibria when the bidders bid uniformly, and that this also holds when valuations are discrete but bids are not. We then go on to consider mixed-strategy Nash equilibria where bidders can bid high in one auction and low in all others. By discretising the bid space, we use smooth fictitious play to compute approximate solutions. Specifically, we find that the results do indeed converge to ϵ-Nash mixed equilibria and, therefore, we are able to locate equilibrium strategies in such complex games where no known solutions previously existed. We then apply this method to compare the auctioneer's revenue for simultaneous second- and first-price auctions. The results show that, for our setting, the revenue-equivalence theorem no longer holds and the first-price auction generates a higher profit for the auctioneer.

262A

Achieving Efficient and Equitable Collaboration among Selfish Agents using Spender-Signed Currency

Geert Jonker, Frank Dignum, John-Jules Meyer

We study collaboration among selfish agents in the tactical airport planning domain. This can be seen as a social exchange scenario, in which the efforts of performing tasks are the resources that are being exchanged. We investigate conditions under which a market mechanism with the use of standard currency leads to efficient and equitable exchange among benevolent agents. We show that, if some agents are selfish, the mechanism can become inequitable and therefore unacceptable. A straight forward penalty rule is not enough to restore equity, as it is attractive to deviate from such a rule. As a solution, we present an novel currency system, under which malicious agents can be punished, resulting in efficient and equitable exchange.

263B

Optimized Algorithms for Multi-Agent Routing

Akihiro Kishimoto, Nathan Sturtevant

Auction methods have been successfully used for coordinating teams of robots in the multi-robot routing problem, a representative domain for multi-agent coordination. Solutions to this problem typically use bids computed using the shortest distance between various locations on a map. But, the cost of this shortest-distance computation has not been considered in previous research.

This paper presents a new auction-based algorithm, FastBid, that works to reduce the computational costs associated with bidding in the multi-robot routing problem. When FastBid requires the shortest distance between two locations it uses an automatically constructed abstraction to compute an accurate heuristic estimate of the actual cost. Two methods are presented to compute this heuristic: one based on A* and the other based on Dijkstra's algorithm. In this way, FastBid can effectively approximate the shortest distances needed for bidding, while preserving the benefits of the auction-based algorithm. We also analyze how small modification in the bidding algorithm can reduce the computational load of the bidding process.

Experiments with simulations of the multi-robot routing problem on large maps demonstrate that the FastBid algorithm not only scales much better than previous approaches, but does so with little or no loss in solution quality.

264C

Evolutionary Dynamics for Designing Multi-Period Auctions

Tomas Klos, Gerrit Jan Van Ahee

Mechanism design (MD) has recently become a very popular approach in the design of distributed systems of autonomous agents. A key assumption required for the application of MD is that agents behave rationally in the mechanism or game, since this provides the predictability of agent behavior required for optimal design of the mechanism. In many cases, however, we are confronted with the intractability both of establishing rational equilibrium behavior, as well as of designing optimal mechanisms even if rational agent behavior can be assumed.

In this paper, we study both sides of the problem simultaneously by designing and analyzing a `meta-game' involving both the designer of the mechanism (game, multi-agent system) and the agents interacting in the system. We use coupled replicator dynamics to investigate equilibrium outcomes in this game. In addition, we present an algorithm for determining the expected payoffs required for our analysis, thus sidestepping the need for extensive simulations as in previous work. Our results show the validity of the algorithm, some interesting conclusions about multi-period auction design, and the general feasibility of our approach.

265A

Understanding How People Design Trading Agents over Time

Efrat Manistersky, Raz Lin, Sarit Kraus

As computerized agents are becoming more and more common, e-commerce becomes a major candidate for incorporation of automated agents. Thus, it is vital to understand how people design agents for online markets and how their design changes over time. This, in turn, will enable better design of agents for these environments. We focus on the design of trading agents for bilateral negotiations with unenforceable agreements. In order to simulate this environment we conducted an experiment with human subjects who were asked to design agents for a resource allocation game. The subjects' agents participated in several tournaments against each other and were given the opportunity to improve their agents based on their performance in previous tournaments. Our results show that, indeed, most subjects modified their agents' strategic behavior with the prospect of improving the performance of their agents, yet their average score significantly decreased throughout the tournaments and became closer to the equilibrium agents' score. In particular, the subjects modified their agents to break more agreements throughout the tournaments. In addition, the subjects increased their means of protection against deceiving agents.

266B

Approximate Bidding Algorithms for a Distributed Combinatorial Auction

Benito Mendoza, Jose Vidal

Distributed allocation and multiagent coordination problems can be solved through combinatorial auctions (CAs). However, most of the existing winner determination algorithms (WDAs) for CAs are centralized. The PAUSE auction is one of a few efforts to release the auctioneer from having to do all the work. The PAUSEBID bidding algorithm generates myopically-optimal bids for agents in a PAUSE auction but its running time is exponential on the number of bids. We present new approximate bidding algorithms that not only run in linear time but also increase the utility of the bidders as result of small decrement in revenue.

267C

Comparing Winner Determination Algorithms for Mixed Multi-Unit Combinatorial Auctions

Brammert Ottens, Ulle Endriss

Mixed multi-unit combinatorial auctions are combinatorial auctions in which the auctioneer and the bidders negotiate over transformations rather than over simple goods. By proposing a transformation a bidder is offering to produce a certain set of output goods after having received the specified input goods. Solving such a mixed auction means choosing a sequence of transformations such that the auctioneer ends up with all the goods desired at the lowest possible cost. This is a generalisation of the winner determination problem in combinatorial auctions and cannot be solved using standard winner determination algorithms. In this paper we analyse the computational complexity of the winner determination problem for mixed auctions and compare the performance of two new algorithms and of the original algorithm proposed for the problem. We also discuss suitable ways of generating test sets for this comparison.

268A

Asynchronous Congestion Games

Michal Penn, Maria Polukarov, Moshe Tennenholtz

We introduce a new class of games, asynchronous congestion games (ACGs). In an ACG, each player has a task that can be carried out by any element of a set of resources, and each resource executes its assigned tasks in a random order. Each player's aim is to minimize his

expected cost which is the sum of two terms – the sum of the fixed costs over the set of his utilized resources and the expected cost of his task execution. The cost of a player's task execution is determined by the earliest time his task is completed, and thus it might be beneficial for him to assign his task to several resources. We prove the existence of pure strategy Nash equilibria in ACGs. Moreover, we present a polynomial time algorithm for finding such an equilibrium in a given ACG.

269B
Beyond quasi-linear utility: strategy/false-name-proof multi-unit auction protocols
Yuko Sakurai, Yasumasa Saito, Atsushi Iwasaki, Makoto Yokoo
This paper introduces strategy/false-name-proof multi-unit auction protocols that can handle non-quasi-linear utilities. One almost universal assumption in auction theory literature is that each bidder has quasi-linear utility, except for some works on budget-constrained bidders. In particular, the celebrated Vickrey-Clarke-Groves (VCG) protocol is strongly believed to critically depend on the quasi-linear assumption and will break down if this assumption does not hold. In this paper, we show that with a simple modification, the VCG can handle non-quasi-linear utilities by sacrificing efficiency to a certain extent. Also, an existing false-name-proof protocol (GM-SMA) can handle non-quasi-linear utilities with the same modification. The basic idea of this modification is that tentative allocation and payments are determined assuming quasi-linear utilities, but each bidder can choose the actual number of units to obtain based on his non-quasi-linear utility. These modified protocols only use the gross utility of each bidder, i.e., his utility for units when his payment is zero. Requiring gross utilities only is an advantage since collecting the entire utility function can be costly. However, determining tentative allocation and payments without considering actual non-quasi-linear utilities can cause significant efficiency loss. To improve efficiency without collecting entire utility functions, we develop a new false-name-proof open ascending auction protocol in which each bidder declares his demand for a series of prices. Our simulation results show that this protocol obtains better social surplus than the modified VCG and GM-SMA when bidders have budget constraints.

270C
Abstractions for Model-Checking Game-theoretic Properties in Combinatorial Auctions
Emmanuel Tadjouddine, Frank Guerin, Wamberto Vasconcelos
We are interested in verifying game-theoretic properties for auction protocols in open agent systems. A property may be that the protocol is robust to collusion or deception or that a given strategy is optimal. Model checking provides an automatic way of carrying out such proofs. However it may suffer from state space explosion for large models. To improve the performance of model checking, abstractions were used along with the SPIN model checker. We considered two case studies: the Vickrey auction and a tractable combinatorial auction. Numerical results showed the limits of relying solely on SPIN. To reduce the state space required by SPIN, two property-preserving abstraction methods were applied: the first is the classical program slicing technique, which removes irrelevant variables with respect to the property; the second replaces large data, possibly infinite values of variables with smaller abstract values. This enabled us to model check the strategy-proofness property of the Vickrey auction for unbounded bid ranges and any number of agents.

271A
Winner Determination in Combinatorial Auctions with Logic-based Bidding Languages
Joel Uckelman, Ulle Endriss
We propose the use of logic-based preference representation languages based on weighted propositional formulas for specifying bids in a combinatorial auction. We then develop several heuristics for a branch-and-bound search algorithm for determining the winning bids in this framework and report on their empirical performance. The logic-based approach is attractive due to its high degree of flexibility in designing a range of different bidding languages within a single conceptual framework.

272B
Towards Agents Participating in Realistic Multi-Unit Sealed-Bid Auctions
Ioannis Vetsikas, Nick Jennings
When autonomous agents decide on their bidding strategies in real world auctions, they have a number of concerns that go beyond the models that are normally analyzed in traditional auction theory. Oftentimes, the agents have budget constraints and the auctions have a reserve price, both of which restrict the bids the agents can place. In addition, their attitude need not be risk-neutral and they may have uncertainty about the value of the goods they are buying. Some of these issues have been examined individually for single-unit sealed-bid auctions. However, here, we work towards extending this analysis to the multi-unit case, and also analyzing the multi-unit sealed-bid auctions in which a combination of these issues are present. In this paper, we present the initial results of this work. More specifically, we present the equilibria that exist in multi-unit sealed-bid auctions, when either the agents can have any risk attitude, or the auction has a reserve price.

Agent Based Simulations and Emergent Behaviour

111C
Shared Focus of Attention for Heterogeneous Agents
Jacob Beal
A network of cooperating agents must be able to reach rough consensus on a set of topics for cooperation. With highly heterogeneous agents, however, incommensurable measures and imprecise translation render ordinary consensus algorithms inappropriate. I present a distributed mechanism for shared focus of attention that begins to address these problems, using an engineered emergence approach inspired by recent results on the dynamics of evolution in systems with spatial extent. Simulation shows that the algorithm converges in time proportional to the diameter of the network and gives a range of reasonable settings for the parameters.

114C
Physical Parameter Optimization in Swarms of Ultra-Low Complexity Agents
Ryan Connaughton, Paul Schermerhorn, Matthias Scheutz
Physical agents (such as wheeled vehicles, UAVs, hover-craft, etc.) with simple control systems are often sensitive

to changes in their physical design and control parameters. As such, it is crucial to evaluate the agent's control systems together with the agent's physical implementation. This can consequently lead to an explosion in the parameter space to be considered. In this paper we investigate the use of swarms of ultra-low complexity agents, and address the issue of finding workable physical agent parameters. We describe a technique for reducing the dimensionality of the search space by performing evaluation tasks that can be used to predict near-optimal parameter values for agents in related multi-agent tasks. We validate our approach on an example task, and demonstrate that this technique can greatly reduce the computational resources required to design a multi-agent system.

117C

An adaptive and customizable feedback system for VR-based training simulators

Maite Lopez-Garate, Alberto Lozano-Rodero, Luis Matey

This paper describes a proposal to build an intelligent feedback selection system for Virtual Reality-based training simulators. The system is aimed at generating multimodal feedback in real-time for advising the students while training with the simulator. Focused on driving tasks, we analyze how to customize the system to exhibit different behaviors. We examine both educational and human factors that have influence on the behavior, so that the instructors can use or refine the behavior they prefer in each training session. The selection process is based on the analysis of the information coming from a diagnostic component that is in charge of determining whether the actions of the student driver are correct or incorrect. The system adapts the feedback to the performance of each student, since the process takes into account whether previous feedbacks were ineffective to prevent that the student repeats the same mistakes. The objective is to emulate the behavior of the instructor. In this way, the feedback system can be helpful for him/her, as while the system decides which the appropriate feedback is, the instructor can focus on other instructional tasks.

118A

An Agent Adaptive Model for Self-Organizing Multi-Agent Systems

Candelaria Sansores, Juan Pavón

Self-organizing multi-agent systems (MAS) use different mechanisms to mimic the adaptation exhibited by complex systems situated in unpredictable, heterogeneous and dynamic environments. These mechanisms allow a collection of agents to spontaneously adapt their behavior towards an optimal organization. Several mechanisms have been proposed based on self-organizing properties according to each discipline, for example, the biology domain has inspired an indirect interaction mechanism for MAS based on stigmergy. This paper presents a self-organization approach that exploits several of these self-organizing properties through an agent adaptive architecture and a reinforcement mechanism. In this approach agents behave according to the principle of rationality, trying to satisfy a set of goals. The behavior of an agent is then guided by these goals and a motivation to pursue those goals. Based on their past experience and a feedback affecting their motivations, agents will select the tasks to execute and the roles to play in a changing environment, in this way adapting dynamically their behavior and that of the overall system. This mechanism was designed and implemented using the INGENIAS methodology for modeling MAS. This methodology was chosen because it is based on a set of meta-models that describe a visual MAS language. The modification of these meta-models provides a visual language which facilitates the specification of the proposed self-organizing mechanism.

Industry and Application Track

91A

MasDISPO_xt: Heat and Sequence Optimisation based on Simulated Trading inside the Supply Chain of Steel Production

Sven Jacobi, David Raber, Klaus Fischer

The production of steel normally constitutes the inception of most Supply Chains in different areas. Steel manufacturing companies are strongly affected by bull whip effects. Due to nondeterministic incoming orders and changes of customer requirements on accepted orders, making the right decision at a certain stage can be the difference between earning or loosing a great turnover. Improving their operational efficiency is required to keep a competitive position on the market. Therefore, flexible planning and scheduling systems are needed to support these processes which are based on considerable amounts of data which can hardly be processed manually. Existing systems are dominated by centralized decision making processes, mostly data driven and often not modeling the business processes they should. MasDISPO xt is an agent-based generic online planning and online scheduling system for monitoring of the complete Supply Chain of Saarstahl AG, a globally respected steel manufacturer. This paper concentrates on the creation and optimisation of heats and sequences as a presetting for the production inside the steelwork.

90C

Argumentation-based agents for eProcurement

Paul-Amaury Matt, Francesca Toni, Thanassis Stournaras, Dimitris Dimitrelos

Procurement is the complete process of obtaining goods and services – from preparation and processing of a requisition through to receipt and approval of the invoice for payment. The support for eProcurement is important for realising agent-based eBusiness applications. This paper proposes the use of argumentation-based agents to support the selection of suppliers for goods and services within the negotiation phase of procurement. Argumentation is used to compare candidate suppliers and identify the one that best meets the buyer's business-specific needs. The use of argumentationbased technology presents important advantages over traditional procurement methods such as competitive bidding, direct negotiation or single-source acquisition in that it can cope with qualitative uncertainty and preferences as well as the construction of contracts. We apply the method to the industrial procurement of an eOrdering system.

93C

Agent-Community-Network-based Business Matching and Collaboration Support System

Tsunenori Mine, Kosaku Kimura, Satoshi Amamiya, Ken'ichi Takahashi, Makoto Amamiya

Business matching and collaboration support systems are useful, in particular for small-and-medium sized companies. Most of them developed so far are based on the server-client architecture and provide their services

onWeb servers. They require special administrative facilities, ask users to upload their data for matching between business needs and seeds, and leave to themselves peer-to-peer communication or negotiation between matched companies. Considering these problems, we have been developing an agent-communitynetwork-based business matching and collaboration support system. Our system requires neither any special administrative facilities nor uploading user data to a special server. It furthermore supports secure peer-to-peer communication between users. It is implemented with multi-agent Kodama framework.

92B

OpCog: An Industrial Development Approach for Cognitive Agent Systems in Military UAV Applications

Kai Reichel, Nico Hochgeschwender, Holger Voos

Future applications of unmanned aerial vehicles (UAVs) especially in military missions require the operation of UAVs with a high level of autonomy. Autonomous UAVs could be developed using agent technologies and therefore this paper investigates such an approach from an industrial perspective. Taking into account time, budget and available knowledge on the industrial side and need for UAV operators to understand the behavior of the autonomous system this paper proposes the application of cognitive agents and a design procedure that supports the transition of the pure operational requirements and functional specification into a cognitive agent system, called Operational driven development approach for Cognitive Systems (OpCog).

95B

Towards a Reliable Air Traffic Control

Minh Nguyen-Duc, Zahia Guessoum, Olivier Marin, Jean-François Perrot, Jean-Pierre Briot, Vu Duong

Since critical socio-technical systems include people interacting with equipments in workplaces, their intrinsic reliability problems have been concerned with both these two “actors”. Air Traffic Control (ATC) is going to be such a system in which controllers use a large number of distributed software tools to provide safety ATC services. The reliability of these services relies on the availability of the various tools. Indeed, a partial failure of a tool in use can have tragic consequences. This paper presents a multi-agent approach to this problem. We propose an agent-based decision-aided system that helps controllers in using their multiple software tools in situations where some tools are not available due to technical incidents. We build and test our system in an ATC simulation environment, thus develop an Agent-Based Simulation (ABS). Experimental work has demonstrated the significance of our system to air traffic controllers.

94A

CAMNEP: Agent-Based Network Intrusion Detection System

Martin *Rehak, Michal Pechoucek, Pavel Celeda, Jiri Novotny, Pavel Minarik*

We present a prototype of agent-based intrusion detection system designed for deployment on high-speed backbone networks. The main contribution of the system is the integration of several anomaly detection techniques by means of collective trust modeling within a group of collaborative detection agents, each featuring a specific detection algorithm. The anomalies are used as an input for the trust modeling. In this stage, each agent determines the flow trustfulness from aggregated anomalies. The aggregation is performed by extended trust models that model the trustfulness of generalized situated identities, represented by a set of observable features. The system is based on traffic statistics in NetFlow format acquired by dedicated hardware-accelerated network cards, and is able to perform a real-time surveillance of the gigabit networks.