



Hyderabad Multi-agent Systems School

4th August -12th August, 2004

Sponsored by
International Foundation for Multiagent Systems
Agents, Theories, Architectures and Languages Workshop
&
Autonomous Agents Steering Committee

Venue



INTERNATIONAL INSTITUTE OF INFORMATION TECHNOLOGY
(formerly Indian Institute of Information Technology) (Deemed University)

Welcome!!!

IIIT, Hyderabad is pleased to host the Multi-agent Systems school sponsored by International Foundation for Multiagent Systems, Agents, Theories, Architectures and Languages Workshop, and Autonomous Agents Steering Committee from August 5th to August 12, 2004 at its campus. The objective of agent school is to initiate, cultivate, and proliferate research in multi agent systems in this part of the world. Multi agent systems have been widely applied to address problems ranging from robotics to web services.

In this Agent School, top researchers in multi-agent systems and their applications will give tutorials on their areas of expertise starting from basic concepts to open research problems. The tutorials are structured such that each researcher can give tutorials on two days with ample time left for discussions and participant / researcher interaction. Projects / Assignments and exercises are scheduled in the evenings for the participants to further grasp the material. This time can also be utilized to collaborate and work on a problem among them or with a researcher.

The participants will be selected based on their qualification and what they want to achieve in agent school. All inquiries regarding this school should be addressed to the local organizer.

We look forward to a highly interactive and energetic Agent School at IIIT, Hyderabad, India.

Kamal Karlapalem,
Local Organizer,
Head, Center for Data Engineering,
International Institute of Information Technology
Hyderabad

About Hyderabad

Hyderabad, the capital of Andhra Pradesh, is a city, which combines both the ancient and the modern. It has the most sophisticated five star hotels, shopping malls, eating places (especially those serving the world famous Hyderabadi biryani, a dish of rice and meat) and entertainment facilities. And it is a maze of markets and tiny houses – a busy, noisy place where narrow ancient lanes meet large vehicle-choked roads. Different in the layout of buildings, the way of life of the people and the very atmosphere, these two cities exist side by side, presenting two very different images. The buildings of present day Hyderabad are a rich mix of Medieval Indian, Saracenic, Mughal and Colonial architecture, a combination of the Hindu and Islamic influences with a hint of the erstwhile British Raj. Hyderabad is highly cosmopolitan, a cauldron of a variety of cultures, not only because of its past but also because of the influx of people from all over the country to its institutions, research and educational, and to its various industries, including many high-profile technological firms. Hyderabad especially draws tourists to its minarets and its pearl bazaar. The gypsy tribes called the 'Lambadas' and 'Banjaras' native to the region are known for their colorful traditions, costumes and handicrafts.

Speakers and Topics

Ed Durfee	10 th & 11 th August	Multiagent Coordination, Planning, and Control
Gal Kaminka	5 th & 6 th August	Multi-Robot Teams, Monitoring Teammates
Sven Koenig	8 th & 9 th August	Decision Making under Time Pressure and Incomplete information
Onn Shehory	7 th & 8 th August	Coalition formation and cooperation/Agent Oriented Software Engineering
Munindar Singh	11 th & 12 th August	Agents and Service-Oriented Computing
Katia Sycara	6 th & 7 th August	Architecting and Deploying Multi Agent Systems/Web Services
Makoto Yokoo	9 th & 10 th August	Multi Agent Systems, Game Theory and Auctions

REGISTRATION***

	By 1st July	By 1st August*	Onsite*
International Participants**			
Full Time Students	USD 200	USD 275	USD 350
Faculty/Post- Doctoral/Government	USD 400	USD 500	USD 600
Industry Participants	USD 600	USD 700	USD 750
Indian Participants**			
Full Time Students	INR 4000	INR 5000	INR 6000
Faculty/Post- Doctoral/Government	INR 8000	INR 9000	INR 10000
Industry Participants	INR 12000	INR 13000	INR 15000

*If there is availability. The number of participants is limited, and registration is dependent on the background of the candidate and his/her interest in multi-agent systems.

**Local accommodation has to be arranged by the participants themselves. IIIT does not guarantee accommodation for the participants in its campus.

***Registration Includes Course materials and all meals.

Interested People should send the following to : agents@iiit.net

- Name,
- Designation,
- Affiliation,
- Email-id,
- Curriculum-vitae,
- Write-up of up to 1000 words explaining how the school will be beneficial to them.

Queries can be posed to :

Kamal Karlapalem,
Organizer, Hyderabad Multi-agent Systems School,
Center for Data Engineering ,
International Institute of Information Technology ,
Gachibowli, Hyderabad-500019,Andhra Pradesh, INDIA.
Email : kamal@iiit.net Fax : + 91-40-23001413
Phone : + 91-40-23001967, 23001969 Extn: 140

Detailed Schedule

August 5th,
Thursday

08:30 – 10:00

10:30 – 12:00

14:00 – 15:30 Multi-Robot Teams -- Gal Kaminka

16:00 – 17:30 Multi-Robot Teams -- Gal Kaminka

18:30 – 20:30 Discussion/Project Time

August 6th,
Friday

08:30 – 10:00 Architecting and Deploying Multiagent Systems
-- Katia Sycara

10:30 – 12:00 Architecting and Deploying Multiagent Systems
-- Katia Sycara

14:00 – 15:30 Monitoring Teammates -- Gal Kaminka

16:00 – 17:30 Monitoring Teammates -- Gal Kaminka

18:30 – 20:30 Discussion/Project Time

August 7th,
Saturday

08:30 – 10:00 Agent-based Web Services -- Katia Sycara

10:30 – 12:00 Agent-based Web Services -- Katia Sycara

14:00 – 15:30 Coalition Formation and Coordination
-- Onn Shehory

16:00 – 17:30 Coalition Formation and Coordination
-- Onn Shehory

18:30 – 20:30 Discussion/Project Time

August 8th,
Sunday

08:30 – 10:00 Decision Making under Time Pressure and
Incomplete Information -- Sven Koenig

10:30 – 12:00 Decision Making under Time Pressure and
Incomplete Information -- Sven Koenig

14:00 – 15:30 Agent Oriented Software Engineering
-- Onn Shehory

16:00 – 17:30 Agent Oriented Software Engineering
-- Onn Shehory

18:30 – 20:30 Discussion/Project Time

August 9th, Monday	08:30 - 10:00	Decision Making under Time Pressure and Incomplete Information -- Sven Koenig
	10:30 - 12:00	Decision Making under Time Pressure and Incomplete Information -- Sven Koenig
	14:00 - 15:30	Multi Agent Systems, Game Theory and Auctions -- Makoto Yokoo
	16:00 - 17:30	Multi Agent Systems, Game Theory and Auctions -- Makoto Yokoo
	18:30 - 20:30	Discussion/Project Time
August 10th, Tuesday	08:30 - 10:00	Multiagent Coordination, Planning, and Control -- Ed Durfee
	10:30 - 12:00	Multiagent Coordination, Planning, and Control -- Ed Durfee
	14:00 - 15:30	Multi Agent Systems, Game Theory and Auctions -- Makoto Yokoo
	16:00 - 17:30	Multi Agent Systems, Game Theory and Auctions -- Makoto Yokoo
	18:30 - 20:30	Discussion/Project Time
August 11th, Wednesday	08:30 - 10:00	Multiagent Coordination, Planning, and Control -- Ed Durfee
	10:30 - 12:00	Multiagent Coordination, Planning, and Control -- Ed Durfee
	14:00 - 15:30	Agents and Service-Oriented Computing -- Munindar Singh
	16:00 - 17:30	Agents and Service-Oriented Computing -- Munindar Singh
	18:30 - 20:30	Project Presentations
August 12th, Thursday	08:30 - 10:00	Project Presentations
	10:30 - 12:00	Project Presentations
	14:00 - 15:30	Agents and Service-Oriented Computing

	-- Munindar Singh
16:00 - 17:30	Agents and Service-Oriented Computing -- Munindar Singh
18:30 - 20:30	Farewell Dinner

ABOUT SPEAKERS

Edmund H. Durfee

Edmund H. Durfee is a Professor at the University of Michigan, where he has been on the faculty for over 15 years. He has conducted research and published extensively in the areas of coordination in multi-agent systems and real-time intelligent control. He is an associate editor of the Journal of Autonomous Agents and Multi-agent Systems, and is a Fellow of IEEE and of AAIL.

Gal Kaminka

Dr. Kaminka is a senior lecturer at the computer science department, Bar Ilan University, and adjunct assistant professor at the computer science department, Carnegie Mellon University. At Bar Ilan University, he heads the MAVERICK research group, which carries out research in robot teamwork, multi-agent systems, and modeling others from observations. He received his Ph.D. in computer science from the University of Southern California (USC). Dr. Kaminka's research activities include building systems and algorithms for behavior recognition, diagnosis of distributed systems, autonomous and operator-controlled multi-robot teams, and agents for computer games.

Sven Koenig

Sven Koenig is an associate professor in the computer science department of the University of Southern California. He received a Ph.D. degree from Carnegie Mellon University for his thesis on "Goal-Directed Acting with Incomplete Information." He also holds M.S. degrees from the University of California at Berkeley and Carnegie Mellon University.

Sven is interested in intelligent systems, especially those that have to operate in large, non-deterministic, non-stationary, or only partially known

domains. His research centers around techniques for decision making (planning and learning) that enable situated agents (including mobile robots or decision-support systems) to act intelligently in their environments and exhibit goal-directed behavior in real-time, even if they have only incomplete knowledge of their environments, imperfect abilities to manipulate them, limited or noisy perception, or insufficient reasoning speed.

Sven has edited one book and published over 70 papers in various areas of artificial intelligence and robotics, including papers in IJCAI, AAAI, AAMAS, AIPS, ECP, ICML, COLT, NIPS, KR, ICRA and IROS. He co-chaired the International Conference on Automated Planning and Scheduling in 2004, the Symposium on Abstraction, Reformulation, and Approximation in 2002, various AAAI and AIPS workshops from 1997 to 2000, as well as the Student Abstract and Poster Sessions at the National Conference on Artificial Intelligence (AAAI) in 1999, 2000 and 2002. He has given over 30 invited talks at research institutions as well as tutorials at the National Conference on Artificial Intelligence, the International Conference on Artificial Intelligence Planning Systems, and the International Conference on Robotics and Automation. He is the recipient of an IBM Faculty Partnership Award, an NSF CAREER award, a Raytheon Faculty Fellowship Award, the Tong Leong Lim Pre-Doctoral Prize and a Fulbright Fellowship. He is proud of the fact that several of his students have won awards for their research.

Onn Shehory

Onn Shehory is a researcher at IBM, Haifa Research Labs, and an adjunct faculty at the Technion, Israel Institute of Technology. He received his Ph.D. in computer science from Bar Ilan University, Israel. Prior to joining IBM Research, Dr. Shehory was a visiting faculty at the Robotics Institute, Carnegie Mellon University, where he conducted research on agents and multi-agent systems. At IBM, he studied agent applications for electronic commerce and contributed to the development of a leading electronic commerce product. In the last decade, he has performed multiple studies of both theory and practice in the fields of distributed artificial intelligence, agents and multi-agent systems, electronic commerce, software engineering, networked storage and autonomous computing. He has a comprehensive knowledge of the leading technologies in these fields, and published dozens of papers in leading journals and conferences, of which two have received best paper awards in recent years. He is on the editorial board of the Journal of Autonomous Agents and Multi-Agent Systems and served as a chair of multiple scientific meetings in the field of agents and electronic commerce

MUNINDAR P. SINGH

Munindar is a full professor in the department of computer science at North Carolina State University. From 1989 through 1995, he was with the Microelectronics and Computer Technology Corporation (MCC). Munindar's research interests include multiagent systems and Web services, where he specifically addresses the challenges of trust, service discovery, and business processes and protocols in large-scale open environments.

Munindar is widely published and has over 150 articles to his name. Munindar's 1994 book *Multiagent Systems*, was published by Springer-Verlag. He coedited *Readings in Agents*, which was published by Morgan Kaufmann in 1998. Munindar is editing the *Practical Handbook of Internet Computing* to be published by CRC Press in 2004 and is coauthoring a new text, tentatively titled, *Service-Oriented Computing*.

Munindar was the editor-in-chief of *IEEE Internet Computing* from 1999 to 2002 and continues to serve on its editorial board. He is also a member of the editorial boards of the *Journal of Autonomous Agents and Multiagent Systems* and the *Journal of Web Semantics*, and serves on the steering committee for the *IEEE Transactions on Mobile Computing*.

Munindar's research has been recognized with awards and sponsorship from the National Science Foundation, DARPA, IBM, Cisco Systems, and Ericsson.

Munindar obtained a B.Tech. in Computer Science and Engineering from the Indian Institute of Technology, Delhi in 1986 and a Ph.D. in Computer Sciences from the University of Texas at Austin in 1993.

Katia Sycara

Prof. Katia Sycara is a Professor in the School of Computer Science at Carnegie Mellon University. She is also the Director of the Laboratory for Semantic Web and Agents Technology. She holds a B.S in Applied Mathematics from Brown University, M.S. in Electrical Engineering from the University of Wisconsin and PhD in Computer Science from Georgia Institute of Technology. She has given numerous invited talks, and has authored or co-authored more than 200 technical papers dealing with Multiagent Systems, Software Agents, Agent Teams, Web Services, the

Semantic Web, Human-Agent Interaction, Negotiation, Case-Based Reasoning and the application of these techniques to e-commerce, crisis action planning, scheduling, manufacturing and financial planning. Prof. Sycara's group has developed the RETSINA multi-agent infrastructure, a toolkit that enables the development of software agents that can dynamically coordinate in open information environments (e.g. battlefield, internet). In addition, Prof. Sycara is one of the contributors to the development of DAML-S/OWL-S, the Darpa-sponsored language for Semantic Web services, as well as matchmaking and brokering software for agent discovery, service integration and semantic interoperation. From 2001-2003 she served as Invited Expert of the W3C (the World Wide Web Consortium) Working Group on Web Services Architecture. She is a member of the Scientific Advisory Board of France Telecom, 2003-2006p;

Prof. Sycara is a Fellow of the American Association for Artificial Intelligence and the recipient of the 2002 ACM/SIGART Agents Research Award. She has served as the Program Chair of the Second International Semantic Web Conference (ISWC 2003), as General Chair of the Second International Conference on Autonomous Agents (Agents 98), as the Chair of the Steering Committee of the Agents Conference (1999-2001), as the Scholarship chair of AAI (1993-1999) and as a member of the AAI Executive Council (1996-99). She is a founding member and member of the Board of Directors of the International Foundation of Multiagent Systems (IFMAS). She is a founding member of the Semantic Web Science Association, and the US co-chair of the Semantic Web Services Initiative. She is a founding Editor-in-Chief of the journal .Autonomous Agents and Multiagent Systems.; an Editor-in-Chief of the Springer Series on Agents; on the Editorial Board of the Kluwer book series on .Multiagent Systems, Artificial Societies and Simulated Organizations.; the Area Editor for AI and Management Science of the journal .Group Decision and Negotiation.. She also serves on the editorial board of the journal .Concurrent Engineering: Research and Applications.. She has served on the editorial board of the .ETAI journal on the Semantic Web. (1998-2001), on the Editorial Board of .IEEE Intelligent Systems and their Applications. (1992-1996), and .AI in Engineering. (1990-1996). She is a member of AAI, the ACM, and Senior Member of IEEE. Prof. Katia Sycara is a Professor in the School of Computer Science at Carnegie Mellon University. She is also the Director of the Laboratory for Semantic Web and Agents Technology. She holds a B.S in Applied Mathematics from Brown University, M.S. in Electrical Engineering from the University of Wisconsin and PhD in Computer Science from Georgia Institute of Technology. She has given numerous invited talks, and has authored or co-authored more than 200 technical papers dealing with Multiagent Systems, Software Agents, Agent Teams, Web Services, the Semantic Web, Human-Agent Interaction, Negotiation, Case-Based Reasoning and the application of these techniques to e-commerce, crisis action planning, scheduling, manufacturing and financial planning. Prof. Sycara's group has developed the RETSINA multi-agent infrastructure, a

toolkit that enables the development of software agents that can dynamically coordinate in open information environments (e.g. battlefield, internet). In addition, Prof. Sycara is one of the contributors to the development of DAML-S/OWL-S, the Darpa-sponsored language for Semantic Web services, as well as matchmaking and brokering software for agent discovery, service integration and semantic interoperation. From 2001-2003 she served as Invited Expert of the W3C (the World Wide Web Consortium) Working Group on Web Services Architecture. She is a member of the Scientific Advisory Board of France Telecom, 2003-2006.

Makoto Yokoo

Makoto Yokoo received the B.E. and M.E. degrees in electrical engineering, in 1984 and 1986, respectively, from the University of Tokyo, Japan, and the Ph.D. degree in information and communication engineering in 1995, from the University of Tokyo, Japan. From 1986 to 2004, he was a research scientist of Nippon Telegraph and Telephone Corporation (NTT). He is currently a Professor of Information Science and Electrical Engineering, Kyushu University. His research interests include multi-agent systems, constraint satisfaction, and mechanism design among self-interested agents. His pioneering work on constraint satisfaction among multiple agents is published in "Distributed Constraint Satisfaction: Foundation of Cooperation in Multi-agent Systems" (Springer, 2001). He received the ACM SIGART Autonomous Agents Research Award in 2004. He served as a program co-chair of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2003). He is an associate editor of Journal of Artificial Intelligence Research. He is on the board of directors of International Foundation for Multiagent Systems (IFMAS) and Japan Society for Software Science and Technology.

ABSTRACTS

Multiagent Coordination, Planning, and Control -- By Ed Durfee

An autonomous agent should coordinate its actions with the actions of other agents to avoid unnecessary conflict and to exploit cooperative opportunities. This requires that agents know enough about what others are doing and plan to do to anticipate how control decisions made now could impact the prospects of future successes. In this tutorial, we will study strategies by which agents can cooperatively solve problems of coordination and control in a practical manner. Topics will include problem representation and decomposition, task and resource allocation, conflict detection and resolution, multi-agent plan improvement, and probabilistic (Markov Decision Process) methods for coordinated distributed control by resource-limited agents.

Multi-Robot Teams -- By Gal Kaminka

Teams of robots are increasingly deployed in real world applications. One of the key challenges in building such teams is to automate the control of teamwork, such that the designer can concentrate her efforts on the taskwork to be done. This tutorial will present methodologies for automating teamwork in teams of behavior-based robots. An instantiated framework and associated algorithms, incorporating many of these methodologies will be presented.

Monitoring Teammates -- By Gal Kaminka

A key challenge in teamwork is to maintain awareness of what others are doing within the context of the team. This is important to maintain agreement among team-members, to allow an operator to monitor the team's progress, and to detect failures. This tutorial will focus on specific challenge areas and present key results and techniques.

Decision Making under Time Pressure and Incomplete Information – By Sven Koenig

Autonomous agents must be able to make good decisions in complex situations that involve a substantial degree of uncertainty, yet find solutions in a timely manner despite a large number of potential contingencies. Unfortunately, decision making in non-deterministic domains is typically time-consuming due to

the large number of contingencies. Thus, autonomous agents need to use decision-making techniques that speed up planning by sacrificing the optimality of the resulting plans, such as agent-centered search and assumption-based planning. In this tutorial, I will give an in-depth overview of such techniques, including algorithms, their analysis using a unifying graph-theoretic framework (including complexity results), and their integration into complete agent architectures. I will then show how these techniques can be used to solve robot-navigation problems, both for single robots and teams of robots. For example, I will show videos of robots that leave trails in the terrain and videos of robots that use auctions for the coordinated exploration of unknown terrain.

Coalition Formation and Cooperation -- By Onn Shehory

Coalition formation is an important cooperation method in multi-agent systems. Within coalitions, agents may be able to jointly perform tasks that they would otherwise be unable to perform, or will perform poorly. To allow agents to form coalitions, one should devise a coalition formation mechanism that includes both a protocol and strategies. Multiple such coalition formation mechanisms were proposed to date. We will review some and discuss underlying assumptions and limitations. We will also discuss recent studies that attempt to relax some of these assumptions thus arrive at automated coalition formation mechanisms suitable for real domains.

Agent Oriented Software Engineering (AOSE) -- By Onn Shehory

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

Architecting and Deploying Multiagent Systems -- By Katia Sycara

Besides being an Artificial Intelligence area of research, multi-agent systems represent a new paradigm in software engineering. In this session, we will present methods for developing multi-agent systems. We will present characteristics of multi-agent systems that engender various architectural requirements. Then, we will present ways of architecting agent systems with emphasis on agent operations in open information environments. We will give examples of multi-agent deployment, with particular focus on our own RETSINA multi-agent infrastructure.

Agent-based Web Services -- By Katia Sycara

Web services are the new wave of software development. Web services are programs that transact on the Web in a distributed fashion. There is a great number of proposed industry standards, the most prominent of those being SOAP, WSDL, UDDI and BPEL. These standards apply to the development of Web services with pre-programmed interfaces and interactions. We claim that research and deployment results from multi agent systems will allow Web services to be more autonomous and flexible. For this to happen an evolution of Web service semantics must take place. In this session, we will discuss the characteristics of Web services, survey current industry standards, present development requirements, and a roadmap that takes Web services from the current state to the future evolution of agent-based Web services.

Multi Agent Systems, Game Theory and Auctions -- By Makoto Yokoo

Game-theory and economics can provide a solid theoretical foundation of autonomous agents and multi-agent systems. Also, due to recent advances in Electronic Commerce, new application fields of autonomous agents and multi-agent technologies, including Internet auctions, have been growing very rapidly. However, some basic concepts of game-theory and economics, such as equilibrium, rationality, uncertainty, etc., are not easily accessible to students with computer science or engineering backgrounds. This lecture gives an introductory tutorial of these basic concepts and their applications from the viewpoint of computer science.

Agents and Service-Oriented Computing – By **Munindar Singh**

Web services have become an important paradigm for information technology architectures and applications. Because of their ubiquity and importance to industry, Web services are a great application area for agent techniques and a great substrate for developing serious agent-based systems. Agents are a natural fit because of their consideration of openness: autonomy, heterogeneity, and dynamism.

This tutorial will present the key concepts, architectures, theories, techniques, and infrastructure of Web services with an emphasis on composition. To this end, it will introduce ontologies, transactions, processes, orchestration, and choreography, and standards for them. It will overlay conventional techniques with key agent technologies---communications, commitments, protocols, contracts, consistency maintenance, and organizations. Attendees will go away with a solid understanding of Web services, the challenges they face, relevant agent techniques, and the opportunities for applying those techniques to address challenges in Web services.

This tutorial is self-contained. It is accessible to Web programmers, advanced developers, and students. It is based on a book titled Service-Oriented Computing authored by Munindar P. Singh and Michael N. Huhns, which will be published by Wiley this summer.

For more information, please see <http://www.csc.ncsu.edu/faculty/mpsingh/tutorials/>

Motivation and Intended Audience

Web services are becoming increasingly important and the basic infrastructure for Web services is becoming quite common. There is an increasing recognition that agents and multi agent systems can be applied in service computing and service computing can facilitate agent-based systems.

This tutorial will address sophisticated approaches support describing, discovering, and engaging Web services, leading up to service composition. Key topics include semantics, transactions, business processes, quality of service, compliance, and trust. The approaches studied are either based on, or complementary to, agents and multi agent systems.

The intended audience includes those with an interest in how agents and multi agent systems relate to service computing. Typical attendees for past tutorials

have been researchers and practitioners from industry and government, advanced developers, graduate and senior undergraduate students, and university faculty.

Detailed Outline

The services metaphor is catching on rapidly for the development of complex Web applications. Because of the heterogeneity and autonomy of web-sites, it is only natural that we model them as independent services. These services can be engaged through well-defined protocols. Protocols, in this sense, replace programming interfaces as an abstraction for programming.

The first generation of the work on Web services has concentrated on the basic infra structural needs, such as directory services, description languages, and invocation standards. The whole point of having Web services is that they be composed into more complex and more valuable services. Present techniques that are popular within the Web community address the challenges of composition only to a limited extent. Some of the most important higher-level abstractions are not studied within the community; instead classical programming techniques are lifted for Web services. These techniques, such as procedure calls, were developed for traditional closed systems.

A lot more can and should be said about Web services, especially when we view them from the perspective of composition. For example, services in general are not invoked but are engaged, meaning that the interactions one has with them are quite unlike method invocations and are better modeled as parts of extended conversations. Similarly, discovering the right service is more than simply looking up a directory with a method signature.

Some of the key techniques for service composition were developed in the areas of databases, distributed computing, artificial intelligence, and multiagent systems. These are generally established bodies of work that can be readily adapted for service composition. Some additional techniques, although inspired by these areas, must be developed from scratch, so as to address the essential openness and scale of Web applications that previous work did not need to address.

Both classes of key techniques should be incorporated into our best practices for service design and composition. In many cases, they can be applied on top of the existing approaches.

This tutorial seeks to discuss the key concepts in service composition. Its intent is to explain the true purposes of service composition, to evaluate existing approaches, and to present existing techniques from other areas that can be

adopted for service composition, and lastly to introduce emerging techniques for addressing challenges that are unique to service composition.

1. Introduction

- Basic concepts of Web services
- Web services architectures and standards
- Motivations and challenges for composition

2. Web Services Architectures and Standards

- Basic concepts
- The services triangle architecture
- WSDL
- SOAP
- UDDI
- Upgrading the services triangle architecture

3. Description: Modeling and Representation

- Conceptual modeling of information and processes
- Ontologies and knowledge sharing
- Relevant standards: RDF, RDFS, and OWL
- Inferencing and tools
- Matchmaking

4. Engagement Basics

- Peer to peer computing
- Messaging
- Distributed transactions
- Business processes
- Relevant standards: WS-Tx, WS-C, BTP
- Evaluation of current architectures and standards

5. Advanced Engagement

- Exception handling
- Relevant standards: BPEL4WS, WSCI, ebXML
- Relaxed transactions
- Monitoring and compliance

6. Collaboration

- Describing compositions
- Agents
- Protocols

Commitments and contracts
Planning
Negotiation
Consistency maintenance
Relevant standards: FIPA, OWL-S, PSL

7. Discovery and Selection

Distributed credentials
Quality of service
Application-level trust
Reputation mechanisms
Referral systems

8. Synthesis

Engineering composed services
Status and trends