

Special Issue on Semantic Extensions to Middleware

Guest Editorial

The Semantic Web is about creating a global machine-processable Web based on innovative mechanisms to publish, access, manage, use, and integrate the large amounts of data available. Despite promising developments at the technological and outreach levels, a number of challenges, most prominently the openness, dynamics, scale and heterogeneity of the (Semantic) Web environment, still hamper the realization of this vision. Moreover, various other trends around the Internet make semantic data become more and more ubiquitous, adding a new dimension to this problem. The Internet of Things connects smallest sensors and devices that need to be coordinated on the basis of shared data and a shared understanding of vertical domains and application scenarios. As a second example, the Internet of Services envisions a Web of billions of service providers and consumers that interact to collaboratively create an immense service economy. Managing sensor data, or provisioning and consuming services requires automation, as the size and the complexity of the respective settings cannot be feasibly mastered manually. Semantics enables automation, and consequently it is expected that not long from now data and service providers and consumers worldwide will interact, integrate and coordinate billions of semantic artifacts of various types. No doubt, there is a need for novel technologies to provide scalable storage and access infrastructures that cope with the heterogeneity, dynamics, openness, complexity and decentralization of the (Semantic) Web.

Middleware technology is traditionally seen to provide the functionality necessary to abstract providers and consumers from syntactic and technical heterogeneities. However, it is clear that for the Semantic Web, solutions such as CORBA are not appropriate. For a truly Web-style global Semantic Web to become reality, it is essential that novel middleware solutions are developed and that further synergies between existing semantic and non-semantic technologies are exploited. In this special issue on ‘Semantic Extensions to Middleware’ we present five approaches that combine promising middleware and semantic technologies research – in integration rather than in current isolation – to address the challenges of the global Semantic Web. The articles tackle the traditional issues of data and process heterogeneity with a clear focus on large-scale, open and distributed environments, i.e., in scenarios that bring data management and service computing to the scale of the originally envisioned global Semantic Web, rather than limiting it to the dimensions of closed enterprise scenarios, or even singular applications.

The works presented in this issue emerged from an international Workshop of Semantic Extensions to Middleware that was held in conjunction with the 2008 OnTheMove Federated Conferences & Workshops in Monterrey, Mexico. The aim of the workshop was to gather researchers, developers and practitioners from various disciplines who seek for innovative techniques and technologies that could add by to the realization of a semantic middleware layer that enables the Semantic Web in the true Web sense. Five selected research contributions that combine and enhance peer-to-peer, tuplespaces or enterprise service bus middleware with semantic technology were selected to elaborate on their work in this issue.

Davoust et al. present a P2P file-sharing system that applies semantics to manage online communities by means of coordinated and shared documents. The communities and documents are interlinked via semantic constructs, and form the basis for a blackboard-style interaction platform. This contribution successfully merges semantic technologies and established P2P infrastructures to create large-scale and dynamically evolvable online structures.

Tolksdorf et al. apply principles of self-organization to store and seek semantic data in decentralized associative storage infrastructures. In this article, the authors combine so-called swarm algorithms with tuplespace-based middleware solutions in order to realize a scalable storage layer for semantic applications. Similarity measures based on the identifiers of RDF resources are used to determine relationships between different pieces of data and to cluster or co-located them for faster discovery.

Guo et al. address the interoperability issues of a semantics-based publish/subscribe system for large-scale networks. Publish-subscribe is an effective approach to improve scalability in terms of communication in distributed systems. This contribution allows different semantic models and mappings to be efficiently and dynamically loaded in order to enable message routing across heterogeneous applications, and thus to foster understanding at much larger scale.

Roa et al. deal with the issue of semantic interoperability at the level of service computing at Web scale. The authors combine established enterprise service bus technology and Semantic Web Services in order to realize a Semantically-enabled Enterprise Service Bus (SESB). The SESB is particularly effective in overcoming large-scale application integration problems by further automating the main Web service life-cycle management tasks and hitch-hiking the communication facilities of service buses.

Besana et al. tackle the service interoperability problem from a coordination rather a communication point of view. A choreography-based approach to service composition is proposed that builds on the shared understanding of contracts between interacting parties. Constraints about the functionalities required are semantically matched against the capabilities of the services in order to realize coordinated activities in distributed settings.

These contributions should provide the reader with insights into the world of semantic systems and initiate further promising endeavors towards semantic extensions to middleware. The ultimate goal of all these efforts is to make the

Semantic Web vision come true at large. We hope that, together with the previously mentioned workshop series at the OnTheMove Federated Conferences & Workshop, this special issue to the Journal of Software will be an important step towards the establishment of a sustainable community of researchers and developers of middleware solutions for the Semantic Web, or semantic enhancements to existing middleware technology.

Guest Editors

Reto Krummenacher (reto.krummenacher@sti2.at)

Semantic Technology Institute, University of Innsbruck, Innsbruck, Austria

Dr. Elena Simperl (elena.simperl@sti2.at)

Semantic Technology Institute. University of Innsbruck, Innsbruck, Austria

Dr. Ronaldo Menezes (rmenezes@cs.fit.edu)

Dept. of Computer Science, Florida Institute of Technology, Melbourne, FL, USA



Reto Krummenacher holds a degree in communication system engineering from the Swiss Federal Institute of Technology in Lausanne (EPFL). Since 2004 he is researcher and project assistant for the Semantic Technology Institute at the University of Innsbruck with a strong focus on semantic middleware in ubiquitous and large scale computing environments. This work was mostly pursued in the context of the European projects TripCom and SOA4All.



Dr. Elena Simperl works as senior researcher and vice-director of the Semantic Technology Institute at the University of Innsbruck. She obtained her PhD in Computer Science (Dr.rer.nat) from the Free University of Berlin in 2007 and holds a diploma in Computer Science from the Technical University of Munich. She is involved in several EU-funded projects such as TripCom, Salero, LarKC, Active, Service Web 3.0 and SOA4All.



Dr. Ronaldo Menezes is associate professor in computer science at Florida Institute of Technology. He received his BSc in computer science in 1992 from the University of Fortaleza, Brazil, and an MSc in computer science in the field of Parallel Computing Models in 1995 from the State University of Campinas, Brazil. In 2000 he was awarded his Ph.D. in computer science at the University of York, UK. He is currently member of the elected group of officials of the ACM Special Interest Group in Applied Computing. He has published more than 40 refereed papers in international events and journals particularly in the fields of Agent Coordination and Bio-Inspired Computing. He is also a program committee member of many conferences and workshops in Swarm Intelligence, Bio-Inspired Computing and Agent Coordination. He is currently doing his sabbatical on Complex Networks and Swarm Intelligence at Northeastern University working with Prof.

Albert-Laszlo Barabasi.