

# TSC – Triple Space Computing

D. Fensel, R. Krummenacher, O. Shafiq, E. Kühn, J. Riemer, Y. Ding, B. Draxler

5 Triple Space Computing (TSC) has been proposed as communication and coordination paradigm based on the convergence of space-  
6 based computing and the Semantic Web. It acts as a global virtual shared space like middleware to enable communication and  
7 coordination of semantic data based on the principle of publish and read. This paper presents an overview of the work in progress under  
8 Austrian FIT-IT funded TSC project (<http://tsc.deri.at>). It presents the evolution of the TSC framework, overall architecture and its usage  
9 by Semantic Web Services.

10 Keywords: communication; coordination; middleware; Semantic Web Services; space-based computing

## 11 TSC – Triple Space Computing.

12 *Triple Space Computing (TSC) ist ein neuartiges Kommunikations- und Koordinationsparadigma, welches aus einer Kombination von*  
13 *„Tuple Spaces“ und Semantic Web-Technologien entstanden ist. Der global zugängliche virtuelle Space stellt eine Middleware zur*  
14 *Verfügung, welche es ermöglicht, semantische Daten via Publizieren und Lesen auszutauschen. In diesem Artikel wird die Entstehung und*  
15 *Entwicklung vom FIT-IT Projekt TSC (<http://tsc.deri.at>) vorgestellt: die Rahmenbedingungen und Datenmodelle, eine Architektur und*  
16 *Anwendungsbeispiele im Zusammenhang mit Semantic Web Services.*

17 *Schlüsselwörter: Kommunikation; Koordination; Middleware; Semantic Web Services; Space-based Computing*

18  
19 Eingegangen am 13. November 2006, angenommen am 23. November 2006  
20 © Springer-Verlag 2007

## 21 1. Introduction

22 Aiming at enhancing the facilities for automated information processing on the Internet, Tim Berners-Lee (inventor of the World  
23 Wide Web and Director of the W3C) brought up the vision of the Semantic Web. Since existing Web technologies around URI, HTTP,  
24 and HTML do not support automated processing of Web content, the aim is to develop technologies that allow describing Web  
25 content in a structured manner; furthermore, semantically defined meta-data shall help to overcome the problem of heterogeneity  
26 within the Internet as an open and distributed system. Ontologies have been identified as the basic building block for the Semantic  
27 Web, as they provide machine-processable, semantic terminology definitions.

28 In conjunction with the idea of the Semantic Web, Web Services are proposed as the technology for automated information processing,  
29 thus combining the benefits of the Web with the strength of component-oriented computation. In fact, Web Services promise  
30 to allow automated interaction and seamless integration of several entities of the Web, thus are considered as the technology for next  
31 generation information systems with special regard to Enterprise Application Integration, B2B technologies, and e-commerce. As  
32 initial Web Service technologies around SOAP, WSDL, and UDDI failed to realize the promise of seamless interoperability, the concept  
33 of Semantic Web Services has been conceived. By adding semantics to Web Service descriptions, intelligent inference-based mechanisms  
34 shall allow automated discovery, composition, and execution of Web Services.

35 Space-based computing has its roots in parallel processing. Linda was developed by David Gelernter in the mid-1980s at Yale University.  
36 Initially presented as a partial language design (Gelernter, 1985), it was then recognized as a novel communication model on its  
37 own and is now referred to as a coordination language for parallel and distributed programming. Coordination provides the infrastruc-  
38 ture for establishing communication and synchronization between

39 activities and for spawning new activities. There are many instantiations or implementations of the Linda model, embedding Linda in  
40 a concrete host language. Examples include C-Linda, Fortran-Linda and Shared-Prolog. Linda allows defining executions of activities or  
41 processes orthogonal to the computation language, i.e. Linda does not care about, how processes do the computation, but only how  
42 these processes are created. The Linda model is a memory model. The Linda memory is called tuple space and consists of logical tuples.  
43 There are two kinds of tuples. Data tuples are passive and contain static data, process tuples or “live tuples” are active and represent  
44 processes under execution. Processes exchange data by writing and reading data tuples to and from the tuple space.

45 In 2003 and 2004 there have been discussions and collaborations involving Tim Berners Lee, Dieter Fensel, Eva Kuehn and Frank  
46 Leymann on the relationships between the Semantic Web, Web Services and space-based computing. Based on that, Dieter Fensel  
47 published a technical report about “Triple Based Computing” presenting the idea of a semantically enabled, space-based communication  
48 and coordination middleware as an infrastructure for the Semantic Web and Semantic Web Services. These ideas have  
49 been adopted for the research project “Triple Space Computing” (TSC) funded by the Forschung, Innovation, Technologie – Informationstechnologie (FIT-IT) research  
50 programme in the programme line of “semantic systems and services”. Triple Space Computing inherits the publication-based communication model  
51 from the space-based computing paradigm and extends it with

52 **Fensel, Dieter, Univ.-Prof. Dipl.-Ing. Dr., Krummenacher, Reto, Dipl.-Ing., Shafiq, Omair,** Bakk. techn., Digital Enterprise Research Institute (DERI), University of Innsbruck, 6020 Innsbruck, Austria (E-mail: dieter.fensel@deri.org); **Kühn, Eva, Ao. Univ.-Prof. Dipl.-Ing. Dr., Riemer, Johannes, Dipl.-Ing.,** Institute of Computer Languages, Department of Computer Science, Vienna University of Technology, 1040 Vienna, Austria; **Ding, Ying, Dr.,** Electronic Web Services (eWS) GmbH, 6020 Innsbruck, Austria; **Draxler, Bernd, Dipl.-Ing.,** Thonhauser Data Engineering (TDE) GmbH, 8700 Leoben, Austria.

1 semantics. Instead of sending messages back and forth among par-  
 2 ticipants as in current message-based technologies, TSC-enabled ap-  
 3 plications will communicate by writing and reading RDF triples in the  
 4 shared space.

## 5 2. Triple Space Computing framework

6 Triple Space Computing (TSC) (Fensel, 2004) implies a number  
 7 of requirements that are not addressed in traditional Linda-like  
 8 (Gelernter, 1985) systems. The requirements are depicted in  
 9 Table 1. They are mainly concerned with enhancing tuple spaces  
 10 with Web and Semantic Web technology like resource identification,  
 11 Semantic Web data models (Resource Description Framework, RDF  
 12 (Klyne, Carroll, 2004)) and query functionality (semantic matching).  
 13 Moreover, TSC is expected to consider the extended scope of in-  
 14 teraction – tuple spaces traditionally serve as communication plat-  
 15 form for process and in-house coordination based on a limited  
 16 number of servers, while TSC aims at a virtually global information  
 17 space. This results in additional requirements on security, reliability  
 18 and scalability.

19 With these requirements in mind, the TSC Framework defines  
 20 the data models, matching algorithms, interaction APIs and security  
 21 models at the convergence of space-based computing/shared  
 22 object spaces and the Semantic Web. The former takes influence  
 23 on the interaction patterns and data matching, while the latter  
 24 determined in particular the RDF-based data modeling approach  
 25 and the storage/query engine installation. In other words TSC  
 26 borrowed from space-based computing its access primitives, trans-  
 27 actional support and eventing/notification mechanism, while the  
 28 Semantic Web provides the RDF triple syntax and semantics with  
 29 the resource identification mechanism (URI) and vocabulary separa-  
 30 tion mechanism (namespaces). Moreover, the RDF query lan-  
 31 guages heavily influenced the definition of the matching mechanism  
 32 through SPARQL (Prud'hommeaux, Seaborne, 2006) and N3QL  
 33 (Berners-Lee, 2004) technology.

34 In the continuation of this section a closer look at the interaction  
 35 API and semantic matching is given. First, however, we explain the  
 36 core data model concepts.

37 The interaction API provides all the primitives defined in Linda. The  
 38 operation names and functionality was, however, mainly influenced  
 39 by more advanced commercial tuple space products like TSpaces  
 40 (Lehmann, McLaughry, Wyckoff, 1999) and JavaSpaces (Freeman,  
 41 Arnold, Hupfer, 1999). In short the API provides operations for  
 42 writing, reading, removing in blocking and non-blocking manners.  
 43 Moreover, convenience methods like update and count were also  
 44 defined. More detailed descriptions of the different operations are  
 45 given in Listing 1. Note that, in order to allow Web-like communica-  
 46 tion, the traditional template-based read and take were enhanced  
 47 with URI-based primitives that allow the extraction of information  
 48 by identifier.

49 In addition to the core API shown in Listing 1 the TSC Framework  
 50 provides APIs for the publish/subscribe extension, the definition of  
 51 mediation rules, the management of spaces (creation, destruction),  
 52 the handling of transaction (commit, rollback) and the definition of  
 53 roles, permissions and users needed for the security framework. A  
 54 user is associated with a particular role, while for every role and  
 55 space the according access permissions can be set.

56 The semantic template matching mechanisms was motivated  
 57 by recent achievement in RDF query languages. A template in Linda  
 58 is a tuple where any of the tuple fields can be replaced by place-  
 59 holders, so-called variables. In TSC templates are defined to be graph  
 60 patterns (detailed definition in (Prud'hommeaux, Seaborne, 2006)).  
 61 The graph patterns are RDF in Notation3 (N3, (Berners-Lee, 2001)),  
 62 where variables can take the place of RDF nodes (cf. Listing 2).  
 63 As graph patterns are at the basis of most RDF query languages  
 64 (in particular SPARQL) the semantic templates can quite easily be  
 65 transformed into queries according to the persistence framework  
 66 (query engine) installed for the Triple Space (cf. Section 3). In that  
 67 way the semantic templates of TSC provide a simple and extensible

**Table 1. Triple Space Computing requirements**

Web-like communication	Application and support of established Web technology like URI for resource identification, stateless exchange of information like supported by HTTP
Publishing mechanism	An interaction model based on the publication of information instead on the exchange of messages
Persistent storage	To ensure decoupling in time and in order to provide the 'publish and read'-paradigm the space must ensure persistency of data
Notification mechanism	For improved coordination and process flow decoupling the installation of a notification mechanism (publish-subscribe paradigm) is required
Search and query	Alignment of Linda-like template matching with Semantic Web query languages in order to provide semantic template matching
Trust and security	Any global information space must ensure confidentiality, integrity, non-repudiation and a trust mechanism to ensure a reasonable service middleware
Reliability	The requirement for persistency implies reliable recovery in case of system failure, as well as transactional support for atomic operations
Versioning	Access logging and tracing of changes is very important when sharing dynamic information in public spaces

**Table 2. TSC data model concepts**

Triple Space	A Triple Space is a uniquely addressable unit of the virtually global information space, i.e., every Triple Space has its own URI. The global space is built by an number of disjoint Triple Spaces (Fig. 1)
Triple	The data model of TSC inherits the syntax (Klyne, Carroll, 2004) and semantics (Hayes, McBride, 2004) of an RDF triple and the same definitions count
Graph	A graph (RDF graph) is defined in (Klyne, Carroll, 2004) as: An <i>RDF graph</i> is a set of RDF triples. Here too, TSC inherits the implied semantics
Named Graph	Named graphs (Carroll et al., 2005) are the fundamental data unit of TSC and as such all communication is based on named graphs. A named graph is a set of triples named by an URI, i.e. the pair (URI name, RDF graph g)

**Listing 1. TSC Interaction API****write (URI ts, Transaction tx, Graph g): URI**

The write operation is used to publish an RDF graph to a triple space identified by the URI ts; the graph name is created by the space upon termination of the write operation and the data stored internally as named graph. Transactional write is supported

**read (URI ts, Transaction tx, Template t): NamedGraph****take (URI ts, Transaction tx, Template t): NamedGraph****query (URI ts, Transaction tx, Template t): Graph**

These three template-based operations are applied to retrieve information from the space. Take has the same semantics as read (retrieval of an entire named graph), however, in a destructive manner. The query primitive on the other hand is used to aggregate all matching triples from the space ts independently of the associated RDF graph, thus it returns a new Graph instead of a whole NamedGraph object. Transactional interaction is supported here, too

**read (URI ts, Transaction tx, URI n): NamedGraph****take (URI ts, Transaction tx, URI n): NamedGraph**

These two operations have the same semantics as their counterparts introduced above. However, they allow retrieving named graphs by use of their name (URI n)

**waitToRead (URI ts, Transaction tx, Template t, TimeOut to): NamedGraph****waitToTake (URI ts, Transaction tx, Template t, TimeOut to): NamedGraph****waitToQuery (URI ts, Transaction tx, Template t, TimeOut to): Graph**

The waitTo-operations (the name was taken from TSpaces) provide blocking versions of the retrieval primitives.

While the previously introduced operations return with NULL in case no data was detected, the blocking versions wait until some match is detected or the timeout runs out. The semantics is otherwise precisely as above

**update (URI ts, Transaction tx, NamedGraph ng): boolean**

Update is on the one hand a convenience method for take and write, and on the other it ensures that graph names are only created by the space. Updates can only be done on graphs that are already known by name to the space. Here too, transactional update is supported

**count (URI ts, Transaction tx, Template t): long**

Count provides the exact same functionality as a loop with counter over a query operation. Note that count only provides an estimate; just as all other primitives the returned set of triples is not ensure to be complete

**Listing 2. Graph pattern-based semantic template**

?s a doap:Project; foaf:member ?o	This graph pattern queries all triples where the subject is of type doap:Project and where the same subject has triples indicating the members
?s ?p ?o. ?o a foaf:Person	This template matches all triples where the object is of type foaf:Person
?s foaf:name ?a; foaf:inbox ?b	This last template matches the triples that contain subjects for which the name and a mailbox (foaf:inbox) are indicated


1 means to match data in a Triple Space, analogue to tuple templates  
2 in Linda.

**3. Triple Space kernel**

4 Like with the Web, the TSC project proposal aimed at building a  
5 Triple Space Computing infrastructure based on the abstract model  
6 called REST (Representational State Transfer) (Fielding, 2000). The  
7 fundamental principle of REST is that resources are identified by URIs  
8 and accessed via a stateless protocol like HTTP in order to transfer  
9 representations, such as HTML or XML documents, of resources over  
10 the network. HTTP provides a minimal set of operations enough to  
11 model any applications domain (Fielding, 2000).

12 Since every representation transfer must be initiated by the client,  
13 and every response must be generated as soon as possible (the  
14 statelessness requirement) there is no way for a server to transmit  
15 any information to a client asynchronously in REST. Furthermore,  
16 there is no direct way to model a peer-to-peer relationship (Khare,  
17 Taylor, 2004) between clients. Finally, HTTP caching based on expira-  
18 tion times for cached requests is not applicable in TSC; where a  
19 server has no pre-knowledge of the lifetimes of named graphs. The  
20 limitations of REST in the context of TSC motivated our approach of  
21 a hybrid architecture called super-peer architecture, which combines  
22 traditional client/server and peer-to-peer architectures. In this archi-  
23 tecture there are three kinds of nodes: servers, heavy clients and  
24 light clients. In the simplest configuration, a particular Triple Space is

realized by a single server, which is accessed by multiple light clients, 25  
for example via HTTP, in order to write and read named graphs and 26  
to receive notifications about graphs of interest. As the number of 27  
light clients increases, the server may become a bottleneck. To over- 28  
come this, additional servers can be deployed to provide additional 29  
access points to a Triple Space for light clients. As a result, a single 30  
Triple Space is be effectively spanned by multiple servers, which use 31  
an inter-server protocol to consistently distribute and collect named 32  
graphs to and from other involved servers. Servers can also be de- 33  
ployed to act as caching proxies in order to improve clients-perceived 34  
access times. The third kind of nodes is heavy clients, which are not 35



**Fig. 1. Definition of the Triple Space**

1 always connected to the system. Like servers they are capable to  
 2 store and replicate Triple Spaces and support users and applications  
 3 to work off-line with their own replicas. While heavy clients can  
 4 join existing Triple Spaces spanned by servers, they are not forced  
 5 to do so.

6 The core functionality of TSC servers and heavy clients is realized  
 7 by a component called Triple Space kernel (TS kernel). Heavy clients  
 8 run in the same address space as the TS kernel, and the TS kernel is  
 9 accessed by its native interface. Light clients use proxies to access  
 10 the TS kernel of a server node transparently over the network. As a  
 11 variation a light client can access a TS kernel via a standardized  
 12 protocol like HTTP, as already mentioned above. In this case a server  
 13 side component, e.g. a servlet, translates the protocol to the native  
 14 TS kernel interface. Figure 2 shows the architecture of the TS kernel.  
 15 Main components of the TS kernel will be briefly described in sub-  
 16 sections below.




Fig. 2. Triple Space kernel architecture

17 **3.1 Mediation engine**

18 Due to diversity in the nature of different communicating partici-  
 19 pants over Triple Space, the possibility of the heterogeneity in the  
 20 data used for communication of different participants may arise and  
 21 make mediation an important issue to be resolved in the Triple Space  
 22 Computing. The Mediation Engine (*Shafiq et al., 2006b*) as part of  
 23 the TS kernel (*Riemer et al., 2007*) is concerned with handling this  
 24 heterogeneity by resolving possibly occurring mismatches among  
 25 different triples. Assume two TSC participants using different data  
 26 models for communication. Then an RDF instance in an RDF schema  
 27 of one TSC participant is needed to be represented in the RDF  
 28 schema of the other TSC participant without altering or loosing the  
 29 semantics. For this reason, a mapping language is needed that  
 30 specifies how to transform the RDF triples according to different  
 31 RDF Schemas of different communicating participants. The media-  
 32 tion rules are to be specified at design time and will be processed by  
 33 a mediation engine at runtime.

34 The TSC mediation engine starts working when users add media-  
 35 tion mapping rules via mediation management interface. Rules  
 36 are defined in the Abstract Mapping Language (AML) (*Scharffe,  
 37 de Bruijn, 2005*) which is independent of any programming lan-  
 38 guage and is able to model complex correspondences that may  
 39 stand between two ontologies. Graphical user interfaces are avail-  
 40 able to define rules in AML. In TSC mediation rules are themselves  
 41 stored in Triple Spaces as RDF graphs. As a result, rules created  
 42 at one server or heavy client can be shared with all other nodes  
 43 spanning a Triple Space. To represent mediation rules in RDF, an  
 44 RDF grounding for AML was defined. A component called mediation

manager implements serialization of rules and helps adding, repla- 45  
 cing and deleting mapping rules. 46

**3.2 Coordination layer** 47

The coordination layer has three responsibilities, (1) local TS opera- 48  
 tions, such as reading and writing named graphs are executed by 49  
 accessing the local data access layer and by consistently propagating 50  
 changes to other involved TS kernels, (2) changes of a space origi- 51  
 nating from other TS kernels are recognized and applied to the local 52  
 data access layer, and (3) remote TS kernels involved to span a 53  
 certain space are discovered automatically in the network. 54

Consistent concurrent access to named graphs is provided via 55  
 transactions. In principle both optimistic and pessimistic transac- 56  
 tions are applicable for TSC; however, they are not exchangeable 57  
 due to differences in their semantics. We decided to support optimistic 58  
 transactions, because they provide a higher degree of concurrency, 59  
 if read operations are more frequent than write operations, which 60  
 results in a higher throughput, because they are free of deadlocks 61  
 without the introduction of additional, semantically sophisticated 62  
 timeout parameters and finally, because they enable a pragmatic 63  
 integration of a data access layer, which itself does not support a 64  
 transaction interface. 65

The prototype implementation of the coordination layer is based 66  
 on the CORSO (Coordinated Shared Objects Spaces) (*Kühn, 1994*) 67  
 middleware. CORSO is a peer-to-peer implementation of a virtual 68  
 shared data space, which allows reading and writing structured, 69  
 shared data objects. It has a built-in distributed transaction manager 70  
 and distributes spaces via an asynchronous, primary-based replica- 71  
 tion protocol. In the TSC prototype, Triple Spaces and named graphs 72  
 are mapped to distributed CORSO data structures. TSC operations 73  
 like reading and writing named graphs are translated to algorithms 74  
 on these CORSO data structures. CORSO further provides a notifica- 75  
 tion mechanism to get informed about changes in the shared space. 76  
 The coordination layer uses CORSO notifications to react on inserted 77  
 or removed named graphs and to asynchronously update the under- 78  
 lying data access layer. The discovery of TS kernels involved in 79  
 spanning a Triple Space is based on the Domain Name System 80  
 (DNS) for wide area networks and on a new protocol based on 81  
 UDP-multicast and CORSO for local area networks. 82

**3.3 Data access layer** 83

Any Triple Space implementation requires a storage and retrieval 84  
 framework to (1) ensure the desired persistency, (2) to support se- 85  
 mantic template matching based on Semantic Web query languages 86  
 and (3) to provide at least a limited amount of reasoning. In order to 87  
 bind arbitrary data stores and query engines to TS kernels we define 88  
 a Data Access Layer (DAL) which defines operations for storing, 89  
 retrieving and deleting RDF graphs. 90

The prototype implementation of the data access layer is based on 91  
 YARS (Yet Another RDF Store) (*Harth, Decker, 2005*), a lightweight 92  
 persistence framework developed in Java at DERI Galway which uses 93  
 optimized indexes for better query performance. Besides the note- 94  
 worthy performance, the fact that the consortium has access to the 95  
 source code and the implementers through DERI Innsbruck, YARS 96  
 has in particular be chosen as it is constructed to store quads or 97  
 contextualized triples instead of plain RDF triples. This allows for 98  
 direct usage of the chosen data model based on named graphs 99  
 (*Carroll et al., 2005*). 100

One of the main tasks of the data access layer is to translate 101  
 templates into N3QL queries for YARS. To keep the Data Access 102  
 API (DAPI) as simple as possible it only defines one operations to 103  
 retrieve data: retrieve(URI ts, Transaction tx, Template t):Graph. The 104  
 TSC API, however, allows a space user to retrieve data either based 105  
 on templates or by use of the graph name. As the DAPI does not 106

1 directly support an interface for URI-based retrieval it is also neces-  
 2 sary to adapt the operation layer in order to transform those  
 3 requests into templates. First, the URI has to be packed into a graph  
 4 pattern template according to (Riemer et al., 2006) as part of the  
 5 Operation Layer processing. The request is then forwarded in form  
 6 of the template to the DAL, where the template is transformed into  
 7 a N3QL query that can be sent to the YARS servlet.

8 **4. Triple Space Computing for Semantic Web Services**

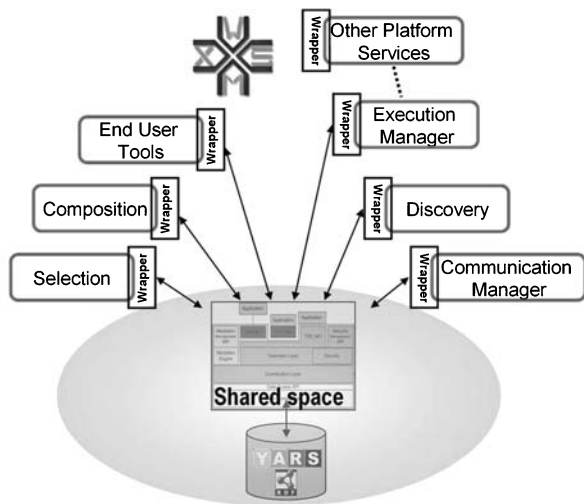
9 Semantic Web Services have been emerged to enable dynamic Web  
 10 Service discovery, composition and execution by using semantic  
 11 descriptions using ontologies as its basis. The Semantic Web Services  
 12 framework has been envisioned (Fensel, Bussler, 2002). It pro-  
 13 vides an ontology, called Web Service Modeling Ontology (WSMO)  
 14 (Roman et al., 2006), a language, called Web Service Modeling  
 15 Language (WSML) (Roman et al., 2006), which provides a formal  
 16 syntax and semantics for WSMO, and an execution environment,  
 17 called Web Service Execution Environment (WSMX) (Roman et al.,  
 18 2006), which is a reference implementation for WSMO, offering  
 19 support for interacting with SWS.

20 The currently used communication paradigm in Semantic Web  
 21 Services (SWS) (Fensel, Bussler, 2002) is synchronous, i.e. users

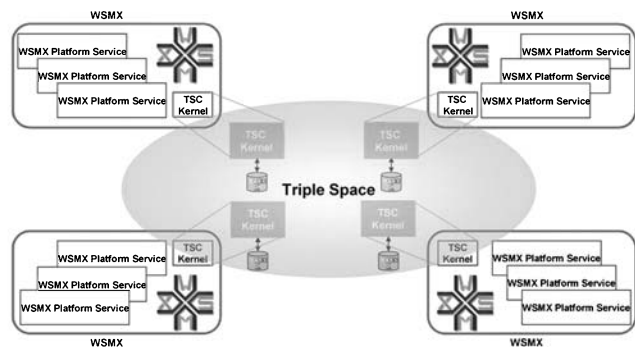
communicate with SWS and SWS communicate with real world  
 Web Services by sending synchronous messages. The problem with  
 synchronous communication is that it requires a quick response as  
 it makes sender halt until the response is received, which is not  
 possible in case of execution process in SWS as it involves heavy  
 processing of semantic descriptions in terms of discovery, selection,  
 composition, mediation, execution. This problem has been over-  
 come by introducing Triple Space Computing as being semantic  
 based asynchronous communication paradigm for communication  
 and coordination of SWS. The Web Services Execution Environment  
 (WSMX) is our reference implementation for SWS in which the Triple  
 Space Computing middleware is being integrated. Using Triple  
 Space Computing in WSMX enables to support greater modulariza-  
 tion, flexibility and decoupling in communication and coordination  
 and to be highly distributed and easily accessible. Multiple TS kernels  
 coordinate with each other to form a virtual space that acts as  
 underline middleware which is used for communication by reading  
 and writing data.

The integration of WSMX and Triple Space Computing has  
 been proposed in four major aspects (Shafiq et al., 2006a): (1)  
 enabling component management in WSMX using Triple Space  
 Computing, (2) allowing external communication grounding in

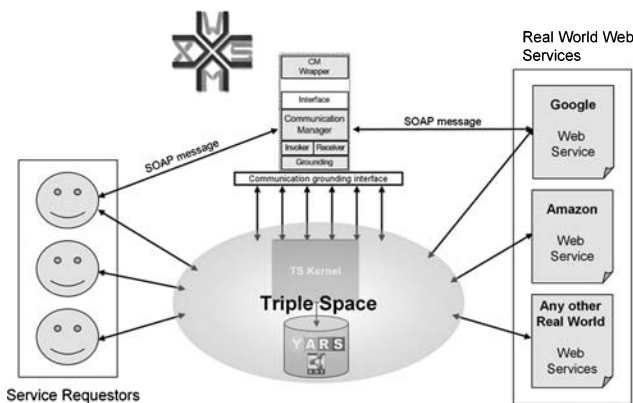
1) Components Management



2) Inter-WSMX Communication and Coordination



3) External Communication Grounding



4) Resource Management




Fig. 3. Triple Space Computing for WSMX

1 WSMX, (3) providing resource management, and (4) enabling communication and coordination between different inter-connected  
2 WSMX systems. Each of the integration aspect is described in the  
3 subsections below. In summary, Triple Space Computing acts as a  
4 middleware for WSMX, Web Services, different other Semantic Web  
5 applications, and users to communicate with each other. Figure 3  
6 shows an initial architecture of each integration aspect.

#### 8 4.1 Component management in WSMX using Triple Space 9 Computing

10 WSMX has a management component that manages the over all  
11 execution of the system by coordinating different components based  
12 on dynamic execution semantics. In this way there has been made a  
13 clear separation between business and management logic in WSMX.  
14 The individual components have clearly defined interfaces and have  
15 component implementation well separated with communication  
16 issues. Each component in WSMX has a wrapper to handle the  
17 communication issues. The WSMX manager and individual compo-  
18 nents wrappers are needed to be interfaced with Triple Space in  
19 order to enable the WSMX manager to manage the components  
20 over Triple Space. The communication between manager and wrap-  
21 pers of the components will be carried out by publishing and  
22 subscribing the data as a set of RDF graphs over triple space.  
23 The wrappers of components that handle communication will be inter-  
24 faced with Triple Space middleware. The WSMX manager has been  
25 designed in such a way that it could distinguish between the data  
26 flows related with the business logic (execution of components based  
27 on the requirements of a concrete operational semantic) and the  
28 data flows related with the management logic (monitoring the com-  
29 ponents, load-balancing, instantiation of threads, etc.).

30 There are two ways for WSMX components to access a TS core,  
31 i.e. heavy clients embed the TS core as a Java package and the  
32 application and TS core run in the same Java Virtual Machine. In  
33 this case CORSO (Kühn, 1994) and YARS (Harth, Decker, 2005)  
34 runtimes need to be deployed together with the heavy client ap-  
35 plication. The second way is to deploy a standalone TS kernel as a  
36 server, which may be accessed by multiple light clients via remoting.  
37 Both scenarios can work. However, we recommend using light  
38 clients in case of communication and coordination within the WSMX  
39 system as in such case it will make the keep the complexity level of  
40 components wrapper and the access of light client embedded in  
41 wrappers will be local to the Triple Space kernel.

#### 42 4.2 Multiple WSMX instances interconnection using Triple 43 Space Computing

44 After enabling WSMX Manager to perform communication and  
45 coordination of components internally, the next step will be to  
46 enable the communication and coordination of different WSMXs  
47 over Triple Space, i.e. forming a cluster of different interconnected  
48 WSMX nodes to support distributed service discovery, selection, com-  
49 position, mediation, invocation, etc. The communication model used  
50 in the current implementation of WSMX is synchronous. Synchronous  
51 communication is beneficial when immediate responses are re-  
52 quired. Since WSMX is dealing with Web service Discovery, Media-  
53 tion and Invocation, immediate responses are usually not available.  
54 In such situations, the synchronous communication will be costly as  
55 it forces the system (component) to remain idle until the response is  
56 available. In order to minimize such overhead imposed by synchrony,  
57 Triple Space can serve as a communication channel between  
58 WSMXs thereby introducing synchronicity between communicating  
59 parties. The Triple Space supports purely asynchronous communica-  
60 tion that optimizes performance as well as communication robustness.

61 The figure above shows the idea of having different WSMX systems  
62 to be interconnected to each other over Triple Space. This will help

the WSMX in providing distributed service discovery, selection, com-  
position, mediation and invocation. There can be the possibility that  
different WSMX systems are running at different location over the  
globe containing different information (i.e. semantic description of  
commercial Web Services, mediation rules, ontologies and goals).  
The service requestor local to a particular WSMX will not be aware  
of other WSMX systems and the data contained by other WSMX  
systems. In this case, it will enable different WSMX systems to be  
aware of each other and to access the data of other WSMXs over  
Triple Space, or redirect the goals to other WSMXs.

#### 4.3 External communication grounding in WSMX using Triple Space Computing

WSMX acts as a semantic middleware between users and real world  
Web Services. Currently, due to existence of message oriented  
communication paradigm, users communicate with WSMX and  
WSMX communicate with Web Services synchronously. The external  
communication manager of WSMX is needed to provide a support  
to communicate over Triple Space. The interfaces for sending and  
receiving external messages by WSMX are needed to provide a  
grounding support to alternatively communicate over Triple Space.  
This needs to be resolved by addressing several issues, i.e. invoker  
component in WSMX is needed to support Web Services Description  
Language (WSDL) and Simple Object Access Protocol (SOAP) com-  
munication binding over Triple Space. The Entry point interfaces will  
be interfaced with Triple Space middleware in order to provide the  
glue between existing Web Services standards and Triple Space  
Computing.

The Communication Manager will be provided with Triple Space  
based grounding support. It will help in providing an additional or  
alternative Triple Space based access interface to access WSMX. It  
will enable Triple Space clients to submit Goals to WSMX via Triple  
Space which will bring the real sense of asynchronous communica-  
tion of Triple Space because normally Goal execution in WSMX  
(performing service discovery, selection, composition, mediation and  
invocation) takes significant amount of time. When the service  
requestors will be able to submit the Goals to WSMX over Triple  
Space, it will not make them hang-up with WSMX until the Goal has  
been executed and will make the communication process of users  
with WSMX more flexible and reliable.

#### 4.4 Resource management in WSMX using Triple Space Computing

WSMX contains different repositories to store ontologies, goals,  
mediators and Web Services descriptions as WSML based files. The  
internal repositories of WSMX are needed to be made optional and  
enable to store the WSML based data as set of RDF named graphs in  
Triple Space Storage. This is mainly concerned with transforming the  
existing representation of data in form of WSML into RDF represen-  
tation. The repository interfaces are needed to be interfaced with  
Triple Space middleware. The Resource Manager in WSMX currently  
manages the persistent storage of data in the repositories. The  
Resource Manager provides a heterogeneous interface for WSMX.  
The component implementing this interface is responsible for storing  
every data item WSMX uses. The WSMO API provides a set of Java  
interfaces that can be used to represent the domain model defined  
by WSMO. WSMO4J (<http://wsmo4j.sourceforge.net>) provides  
both the API itself and a reference implementation. Currently WSMX  
defines interfaces for six repositories. Four of these repositories  
correspond to the top level concept of WSMO, i.e. Web Services,  
ontologies, goals, and mediators. The fifth repository is for non-  
WSMO data items e.g. events and messages. Finally, the sixth repo-  
sitory stores WSDL documents used to ground WSMO service  
descriptions to SOAP.

1 The storage of WSMO top level entities on Triple Space will help in  
 2 enhancing and fastening the process access of the data items after-  
 3 wards. For instance, in the current discovery mechanism of WSMX,  
 4 the WSML reasoners have to reason on each and every Web Service  
 5 description available in the local repositories which takes significant  
 6 amount of time. When the Web Services descriptions will be stored  
 7 over Triple Space, the template matching based simpler reasoning  
 8 will be used as a first step in order to filter-out the most relevant and  
 9 possibly required Web Service descriptions. The filtered Web Services  
 10 descriptions based on template based matching over Triple Space  
 11 are retrieved and converted back to WSML to be used by WSML  
 12 reasoners. It makes the process of discovery simpler and faster by  
 13 performing reasoning operations only on relevant Web Service  
 14 descriptions rather than all.

## 15 5. Conclusions

16 In this paper we provide an overview of overall work in-progress in  
 17 Triple Space Computing (TSC) project funded by the Austrian Gov-  
 18 ernment under the program FIT-IT Semantic System. In the project  
 19 we are building the Triple Space Computing as a communication  
 20 and coordination framework for semantic technologies. In this paper  
 21 we presented the background of TSC, introduction, state-of-the-art,  
 22 TSC framework, data and interaction model, TSC architecture and  
 23 TSC integration with Semantic Web Services. The project has en-  
 24 tered into its final phase where theoretical work has been completed  
 25 and currently prototypes are under development which will be fol-  
 26 lowed by extensive evaluation.

## 27 Acknowledgements

28 The work is funded by the Austrian Government under the program  
 29 FIT-IT (Forschung, Innovation, Technologie – Informationstechnolo-  
 30 gie), project 809249 TSC – Triple Space Computing.

## 31 References

- 32 Berners-Lee, T. (2001): Notation 3: An RDF language for the Semantic Web. W3C Design  
 33 Issues, November 2001, <http://www.w3.org/DesignIssues/Notation3.html>.  
 34 Berners-Lee, T. (2004): N3QL – RDF Data Query Language. W3C Design Issues, July 2004,  
 35 <http://www.w3.org/DesignIssues/N3QL.html>.  
 36 Bussler, C. (2005): A minimal triple space computing architecture. Proc. of the 2nd WSMO  
 37 Implementation Workshop (WIW 2005), Innsbruck, Austria, June 2005.  
 38 Carroll, J. J., Bizer, Ch., Hayes, P., Stickler, P. (2005): Named graphs. Journal of Web  
 39 Semantics 3 (4).  
 40 Fensel, D. (2004): Triple-space computing: Semantic web services based on persistent  
 41 publication of information. Proc. of the IFIP Int. Conf. on Intelligence in Communication  
 42 Systems, INTELLCOMM 2004, Bangkok, Thailand, November 2004.

97

## 98 The Authors



### Dieter Fensel

is Director of Digital Enterprise Research Institute (DERI) at the Leopold Franzens University of Innsbruck, Austria. Prior to this he was Director of Digital Enterprise Research Institute (DERI) at National University Ireland at Galway, Ireland. Moreover, he has held positions at University of Karlsruhe (AIFB), University of Amsterdam (UvA), and the Vrije Universiteit Amsterdam (VU). His current research interests are around the usage of semantics in 21st century computer science. He received his PhD from University of Karlsruhe in 1994.

- Fensel, D., Bussler, C. (2002): The Web Service Modeling Framework WSMF. Electronic  
 Commerce Research and Applications 1 (2). 43  
 Fielding, R. T. (2000): Architectural styles and the design of network-based software  
 architectures. PhD Thesis, University of California, Irvine. 44  
 Freeman, E., Arnold, K., Hupfer, S. (1999): JavaSpaces: principles, patterns, and practice.  
 The Jini Technology Series, Addison-Wesley Longman. 45  
 Gelernter, D. (1985): Generative communication in Linda. ACM Transactions in Prog.  
 Languages and Systems (TOPLAS) 7(1). 46  
 Harth, A., Decker, S. (2005): Optimized index structures for querying RDF from the Web.  
 Proc. of the 3rd Latin American Web Congress, Buenos Aires, Argentina,  
 October/November 2005. 47  
 Hayes, P., McBride, B. (eds.) (2004): RDF Semantics. W3C Recommendation, February  
 2004, <http://www.w3.org/TR/rdf-nt/>. 48  
 Khare, R., Taylor, R. N. (2004): Extending the Representational State Transfer (REST)  
 Architectural style for decentralized systems. Proc. of the Int. Conf. on Software  
 Engineering (ICSE), Edinburgh, Scotland, May 2004. 49  
 Klyne, G., Carroll, J. J. (eds.) (2004): Resource Description Framework (RDF):  
 concepts and abstract syntax. W3C Recommendation, February 2004,  
<http://www.w3.org/TR/rdf-concepts/>. 50  
 Krummenacher, R., Hepp, M., Polleres, A., Bussler, C., Fensel, D. (2005): WWW or What Is  
 Wrong with Web Services. Proc. of the 2005 IEEE European Conf. on Web Services  
 (ECOWS 2005), Växjö, Sweden, November 2005. 51  
 Kühn, E. (1994): Fault-tolerance for communicating multibase transactions. Proc. of the  
 27th Hawaii Int. Conf. on System Sciences (HICSS), Wailea, Maui, Hawaii, January 1994. 52  
 Kühn, E., Beinhart, M., Murth, M. (2005): Improving data quality of mobile Internet  
 applications with an extensible virtual shared memory approach. Proc. of Iadis  
 www/Internet 2005 Conference, Lisbon, December 2005. 53  
 Lehman, T. J., McLaughry, St. W., Wyckoff, P. (1999): TSpaces: The Next Wave. Proc. of the  
 Hawaii Int. Conf. on System Sciences (HICSS-32), Wailea, Maui, Hawaii, January 1999. 54  
 Prud'hommeaux, E., Seaborne, A. (eds.) (2006): SPARQL Query Language for RDF. W3C  
 Candidate Recommendation, April 2006, <http://www.w3.org/TR/rdf-sparql-query/>. 55  
 Riemer, J., Martin-Requerda, F., Ding, Y., Murth, M., Sapkota, B., Krummenacher, R.,  
 Shafiq, O., Fensel, D., Kuehn, E. (2006): Triple Space Computing: adding semantics to  
 space-based computing. Proc. 1st Asian Semantic Web Conf., Beijing, China,  
 September 2006. 56  
 Roman, D., de Bruijn, J., Mocan, A., Lausen, H., Bussler, C., Fensel, D. (2006): WWW –  
 WSMO, WSML, and WSMX in a nutshell. Proc. of the 1st Asian Semantic Web Conf.  
 (ASWC 2006), Beijing, China, September 2006. 57  
 Sapkota, B., Kilgarriff, E., Bussler, C. (2006): Role of Triple Space Computing in Semantic  
 Web Services. Proc. of the 8th Asia Pacific Web Conf. (APWEB 2006), Harbin, China,  
 January 2006. 58  
 Scharffe, F., de Bruijn, J. (2005): A language to specify mappings between ontologies.  
 Proc. of 1st Int. Conf. on signal-image technology and Internet-based systems  
 (SITIS 2005), Yaounde, Cameroon, November 2005. 59  
 Shafiq, O., Krummenacher, R., Martin-Requerda, F., Ding, Y., Fensel, D. (2006a): Triple  
 Space Computing Middleware for Semantic Web Services. Proc. of 2006 Middleware  
 for Web Services (MWS 2006) Workshop at the 10th Int. IEEE Enterprise Computing  
 Conf. (EDOC 2006) in Hong Kong, October 2006. 60  
 Shafiq, O., Scharffe, F., Krummenacher, R., Ding, Y., Fensel, D. (2006b): Data mediation  
 support for Triple Space Computing. Proc. of the 2nd IEEE Int. Conf. on Collaborative  
 Computing (CollaborateCom 2006), Atlanta, GA, USA, November 2006. 61  
 WSMO4: an API and a reference implementation for building Semantic Web Services  
 applications compliant with the Web Service Modeling Ontology (WSMO),  
<http://wsmo4.sourceforge.net/>. 62  
 63  
 64  
 65  
 66  
 67  
 68  
 69  
 70  
 71  
 72  
 73  
 74  
 75  
 76  
 77  
 78  
 79  
 80  
 81  
 82  
 83  
 84  
 85  
 86  
 87  
 88  
 89  
 90  
 91  
 92  
 93  
 94  
 95  
 96



### Reto Krummenacher

holds a degree in communication system engineering from the Swiss Federal Institute of Technology in Lausanne, Switzerland (EPFL). He concluded his master studies in 2004 with a thesis on prefetching for location-aware mobile information systems which was conducted at Fraunhofer's Integrated Publication and Information Systems Institute IPSI in Darmstadt, Germany. Since late 2004 he is, as PhD student, part of

the Ubiquitous Services Working Group at the Digital Enterprise Research Institute of the University of Innsbruck (DERI Innsbruck) with a strong focus on Triple Space Computing in ubiquitous computing environments. 123  
 124  
 125  
 126



**Omaid Shafiq**

is working as junior researcher at the Digital Enterprise Research Institute (DERI), University of Innsbruck, Austria. Prior to this he was working as research associate in the Semantic Grid and Multi Agent Systems Research Group in the Information Technology department of National University of Sciences and Technology (NUST), Rawalpindi/Islamabad, Pakistan. He received his integrated degree of BS in Com-

puting with honors and Rector's Gold Medal award in June 2005 from the same university. He was also serving as technical leader of joint research initiative between NUST Pakistan, Communication Technologies (Comtec), Sendai, Japan and University of Arkansas at Little Rock (UALR), AR, USA till August 2005. His current research interests focus on Semantic Web, Semantic Web Services, Semantic Grid, Peer to Peer Computing and Distributed Systems Management.



**Eva Kühn**

holds the titles graduated engineer of computer sciences (Dipl.-Ing.), Ph.D. (Dr. techn.) and Venia Docendi (Univ.-Doz.) from the Vienna University of Technology, Vienna, Austria. She received the Heinz-Zemanek research award for the her Ph.D. work on "Multi Database Systems" and a Kurt-Gödel Research Grant from the Austrian Government for a sabbatical at the Indiana Center for Databases at Purdue University,

Indiana, USA. Dr. Kühn is employed as A.o. Univ.-Prof. at the Vienna University of Technology, Institute of Computer Languages. Her current teaching topics are virtual shared memory systems and coordination tools. Since 2006 Dr. Kühn is a member of the Senate of the Christian Doppler Forschungsgesellschaft.



**Johannes Riemer**

studied computer science at the Vienna University of Technology and received a degree in Computer Science (Diploma) in 2003. During his studies his special interests have been on distributed systems, real-time systems and object oriented languages. Parallel to university he has been working in industry since 1999, where he gained experience in design and development of distributed applications and virtual shared

memory middleware. Since 2005 he is a PhD student at the Vienna University of Technology with a research focus on space-based computing, virtual shared memory, triple space computing, semantic web and service oriented architecture.

47  
48  
49  
50



**Ying Ding**

is senior consultant in Electronic Web Services GmbH. Prior to this she was a researcher at the Division of Mathematics and Computer Science at the Free University of Amsterdam. She completed her Ph.D. in School of Applied Science, Nanyang Technological University, Singapore. She has been involved in various European-Union funded projects and served as consultant in many projects between universities and

the companies. Her current interest areas include semantic web, web services, semantic information retrieval, knowledge and content management, e-commerce and commercial application of semantic web technology.

62  
63  
64  
65



**Bernhard Draxler**

is working as software developer and designer for data acquisition software based on java/C++ for windows and Embedded Linux with SQL-databases at Thonhauser Data Engineering GmbH, Leoben, Austria. He received his diploma in 1997 on Vienna University of Technology. Prior to this current position, he worked on virtual shared memory system for windows and Unix Linux at Tecco software development CORP, at

IBM Austria on development of workflow management systems and at Vienna software publishing team leader development of an Office package for OS/2.

77  
78  
79  
80