Artemis: Deploying Semantically Enriched Web Services in the Healthcare Domain ⋆

A. Dogac, G. Laleci, S. Kirbas, Y. Kabak, S. Sinir, A. Yildiz, Y. Gurcan

Software Research and Development Center

Middle East Technical University (METU)

06531 Ankara Turkiye

email: asuman@srdc.metu.edu.tr

Tel: +903122105598, Fax: +903122101004

Abstract

An essential element in defining the semantics of Web services is the domain knowledge. Medical informatics is one of the few domains to have considerable domain knowledge exposed through standards. These standards offer significant value in terms of expressing the semantics of Web services in the healthcare domain.

In this paper, we describe the architecture of the Artemis project, which exploits ontologies based on the domain knowledge exposed by the healthcare information standards through standard bodies like HL7, CEN TC251, ISO TC215 and GEHR. We use these standards for two purposes: first to describe the Web service functionality semantics, that is, the meaning associated with what a Web service does and secondly to describe the meaning associated with the messages or documents exchanged through Web services.

Artemis Web service architecture uses ontologies to describe semantics but it does not propose globally agreed ontologies; rather healthcare institutes reconcile their semantic differences through a mediator component. The mediator component uses ontologies based on prominent healthcare standards as references to facilitate semantic mediation among involved institutes. Mediators have a P2P communication architecture to provide scalability and to facilitate the discovery of other mediators.

Key words: Medical Information Systems, Web Services, Semantic Web, P2P Technologies, Electronic Healthcare Records (EHR), Interoperability

1 Introduction

Most of the health information systems today are proprietary and often only serve one specific department within a healthcare institute resulting in difficult interoperability problems. To complicate the matters worse, a patient's health information may be spread out over a number of different institutes which do not interoperate. This makes it very difficult for clinicians to capture a complete clinical history of a patient.

On the other hand, the Web services model provides the healthcare industry with an ideal platform to overcome the difficult interoperability problems. Web services are designed to wrap and expose existing resources and provide interoperability among diverse applications.

Introducing Web services to the healthcare domain brings many advantages:

- It becomes possible to provide the interoperability of medical information systems through standardizing the access to data through WSDL [47] and

⋆ This work is supported by the European Commission through IST-1-002103-STP Artemis project and in part by the Scientific and Technical Research Council of Turkey, Project No: EEEAG 104E013

SOAP [44] rather than standardizing documentation of electronic health records.

- Medical information systems suffer from proliferation of standards to represent the same data. Web services allow for seamless integration of disparate applications representing different and, at times, competing standards.
- Web services will extend the healthcare enterprises by making their own services available to others.
- Web services will extend the life of the existing software by exposing previously proprietary functions as Web services.

However it has been generally agreed that Web services offer limited use unless their semantics are properly described and exploited [34–36,38].

Generic service semantics, that is, semantics applicable to all services such as constraints on temporal and spatial availability and service quality, can be defined through DAML-S [8] (later OWL-S) upper ontology. However some other properties and meaning to be attached to the Web services depend on the application domain. The domain specific semantics is necessary for the Web services in the following respects:

- For describing service functionality semantics: In order to facilitate the discovery of the Web services, there is a need for semantics to describe what the service does, which is the service functionality semantics in the domain. For example, in the healthcare domain, when a user is looking for a service to admit a patient to a hospital, he should be able to locate such a service through its meaning, independent of what the service is called and in which language. Note that WSDL [47] does not provide this information.
- For describing service message semantics: When invoking a Web service,

there is also a need to know the meaning associated with the messages or documents exchanged through the Web service. In other words, service functionality semantics may suffice only when all the Web services use the same message standards. For example, a "GetClinicalInformation" Web service may include the messages to pass information on diagnosis, allergies, encounters and observation results about a patient. Each part of the message must either conform to a single EHR standard or should be semantically annotated to make sense at the receiving end.

Hence an essential element in defining the semantic of Web services is the domain knowledge. The healthcare information standards through standard bodies like HL7 [23], CEN TC251 [7] , ISO TC215 [27] and GEHR [21] expose considerable domain knowledge through classifications, methodologies, terminologies, and controlled vocabularies. Although this domain knowledge is useful, there are more benefits to be gained by expressing such knowledge through formal ontology languages like Web Ontology Language (OWL) [46]:

- An ontology language is machine processable since it conforms to a formal, well-defined syntax. A description given in an ontology language can be automatically processed to obtain the metadata. For example, a description in OWL can be parsed into the classes, properties and corresponding values, even when an application knows only the OWL syntax and has no understanding of a particular domain specific ontology. However to interpret the metadata automatically, its meaning (semantics) must also be given through domain specific ontologies. In this way, any program having a prior knowledge of the syntax and semantics of the ontology, can parse the description, extract meta-data and interpret it since the syntax and the semantics is already known by the application using it.

- An ontology describes *consensual knowledge*, that is, it describes meaning which has been accepted by a group not by a single individual; in other words, it provides a common vocabulary for those who have agreed to use it. Hence when we annotate a Web service with a node in an ontology it inherits the well-defined, shared meaning attributed to that node. For example, when a Web service instance, say "HastaKabul" is annotated with the "AdmitPatient" node of a medical ontology, its operational meaning becomes clear that this service can be used in admitting patients to a hospital.
- An ontology provides the ability to define relationships among classes, properties and instances which can then be used for reasoning. For example, assume that each patient has a unique patient identifier which can be expressed through the OWL FunctionalProperty. Then given two patients with different names (probably misspelled during administration) but having the same patient identifier (which is guaranteed to be unique by the system), an OWL reasoner can deduce that these two patients are in fact the same.

In this paper, we present the design and implementation of a semantically enriched Web service based interoperability platform for the healthcare domain which is being developed within the scope of the Artemis project [2]. The domain knowledge exposed by prominent healthcare standards are organized into ontologies and these ontologies are used in associating both "service functionality" and "service message" semantics with the Web services. It should be noted that the ontologies we are proposing are just to facilitate ontology mediation. In other words, we do not find it realistic to expect healthcare institutes to conform to one global ontology. The Artemis architecture allows the healthcare institutes to develop their own ontologies. However, when these ontologies are based on standards developed by the healthcare standardization bodies like CEN TC251, ISO TC215, GEHR or HL7, we show that ontology mappings are facilitated to a great extend through semantic mediation. The mediator architecture in Artemis is based on a peer-to-peer infrastructure to provide scalability and to facilitate the discovery of other mediators.

The paper is organized as follows: In Section 2, we describe how the semantics exposed by the healthcare standards can be taken advantage of in developing a Web service technology framework for healthcare domain. We introduce the *Service Functionality Ontology*, the *Service Message Ontology* and show how to use a MAFRA [32] based ontology mapping mechanism. Complex service composition from elementary services using the semantics and semantic aggregation operators proposed are also discussed in this section. In Section 3, we present the system architecture and the Artemis mediator component. The implementation status of the system is also given in this section. Section 4 describes the related work. Finally Section 5 summarizes the contributions of the paper.

2 Exploiting Web Service Technology in Healthcare Informatics

Medicine is one of the few domains to have extensive domain knowledge defined through standards. Some of the domain knowledge exists in "controlled vocabularies", or "terminologies". Some vocabularies are rich semantic nets, such as SNOMED-CT [43] while others such as ICD-10 (International Statistical Classification of Diseases and Related Health Problems) [25] are little more than lexicons of terms. However, in addition to such vocabularies and taxonomies, there are standards that expose the business logic in the healthcare domain such as HL7 [23]. Electronic Healthcare Record based standards such as CEN TC251 [7], ISO TC215 [27] and GEHR [21], on the other hand, define and classify clinical concepts that make up the patient records. Such standards offer significant value in developing ontologies to express the semantics of Web services.

The semantics is necessary in medical Web services in the following respects:

- First, as mentioned in the Introduction section, in order to facilitate the discovery of the Web services, there is a need for an ontology to describe service functionality in the healthcare domain.
- Describing the service functionality semantics is not enough; in real life medical information services, there can be quite complex service parameters and therefore both the semantics and the structure of the message parameters are also necessary to decipher them at the receiving end.
- As already noted, it is not realistic to expect global ontologies; rather it is possible to have more than one ontology to express the similar concepts. This is especially true for the medical information systems: the EHR based standards use different terminologies for similar concepts.

  Moreover, given these standards, it is also not realistic to ignore all these efforts and develop brand new ontologies. Therefore, it is reasonable to expect healthcare institutes to develop or use their own ontologies based on the concepts provided by the existing healthcare information standards.

  On the other hand, it is possible to specify the mappings between ontologies based on existing standards. Such mappings make it possible to facilitate the mediation between healthcare institutes' own ontologies as long as they make use of ontologies based on these standards.

Furthermore, the semantic constructs developed must be integrated with the service registries which provide the basic mechanisms for service discovery.

There are basically two different healthcare standardization approaches: The first approach is message based such as HL7 [23]; the other is Electronic Health Care Record (EHR) based such as CEN ENV 13606 [7], and GEHR [21]. In the following sections, we describe how these standards can be exploited in developing semantic based healthcare Web services.

2.1 Healthcare Information Standards and Web services

The primary goal of HL7 is to provide standards for the exchange of data among healthcare computer applications. The standard is developed with the assumption that an event in the healthcare world, called the *trigger event*, causes exchange of messages between a pair of applications. When an event occurs in an HL7 compliant system, an HL7 message is prepared by collecting the necessary data from the underlying systems and it is passed to the requestor, usually as an EDI message. For example, a trigger event can occur when a patient is admitted and this may cause the data about that patient to be collected and sent to a number of other systems.

Since HL7 defines message based events, one might think that these events can directly be mapped into Web services. However, this may result in several inefficiencies. The input and output messages defined for HL7 events are usually very complex containing innumerous segments of different types and optionality. Furthermore, all the semantics about the business logic and the document structure are hard coded in the message. This implies that, the party invoking the Web service must be HL7 compliant so that the content of the output parameter(s) returned by the service makes sense.

RQC Request Clinical Information	
MSH	Message Header
QRD	Query Definition
[QRF]	Query Filter
{	
PRD	Provider Data
[{ CTD }]	Contact Data
}	
PID	Patient Identification
[{ NK1 }]	Next of Kin/Associated Parties
[{ GT1 }]	Guarantor
[{ NTE }]	Notes and Comments

RCI Return Clinical Information	
MSH	Message Header
MSA	Message Acknowledgment
[QRF]	Query Filter
{	
PRD	Provider Data
[{ CTD }]	Contact Data
}	
PID	Patient Identification
[{ DG1 }]	Diagnosis
[{ DRG }]	Diagnosis Related Group
[{ AL1 }]	Allergy Information
[	
{	
OBR	Observation Request
[{ NTE }]	Notes and Comments
[	
{	
OBX	Observation Result
[{ NTE }]	Notes and Comments
}	
]	
}	
]	
[{ NTE }]	Notes and Comments

Fig. 1. The Structures of the RQC/RCI messages for the HL7 event I05

Note further that some of the information contained in an HL7 message may be coming from different systems either proprietary or complying to different standards. For example, the event I05 in HL7 is used to pass the clinical patient information given patient identification information. Clinical information refers to the data contained in a patient record such as problem lists, lab results, current medications, family history, etc. [24]. The input and output messages of I05 are shown in Figure 1. All or some of this data may be coming from different systems that do not interoperate. This in turn, creates the need to retrieve these partial results probably through finer granularity Web services. Hence, in Web services terminology, HL7 events correspond to "Composite services", whereas more elementary services are needed. Deciding upon the "elementary" service granularity is important since this affects the service reusability and interoperability with other healthcare standards.

In order to define the granularity of Web services, we refer to Electronic Healthcare Record (EHR) based standards from major standard bodies like CEN and GEHR. These standards define metadata about EHR through "meaningful components".

When a Web service is designed to retrieve a fine granularity “meaningful component” of an EHR, it can be semantically annotated as such. In other words, we propose to annotate the semantics of fine granularity Web services through the semantics of the messages that they carry. In this way, a healthcare institute, say Hospital A, conforming to HL7 can annotate its Web services conforming to an HL7 compliant ontology; on the other hand, another healthcare institute, say Hospital B conforming to CEN can annotate its Web services with a CEN compliant ontology. Note that, the mapping between these ontologies needs to be available at the Artemis Mediator for these two institutes to understand each other.

This provides the following benefits:

- The semantics of Web services can be mapped between different EHR standards to achieve interoperability. For example, a Web service retrieving “Allergy Information” can be semantically annotated as “AL1” in HL7. Then a CEN’s ENV 13606 compliant system can understand the semantics of this service through an ontology mapping indicating that “AL1” in HL7 corresponds to “DF03” in a CEN’s ENV 13606 compliant system where DF03 denotes the Allergy Information in CEN’s ENV 13606.
- Web service reusability is improved; a Web service can not only be invoked by other applications which need only that piece of data but also be used as a component of a larger composite service. For example a service retrieving “Allergy Information” can be a part of a composite service retrieving the whole clinical information about a patient.

As a summary, there is a need for a *Service Functionality Ontology* to classify coarse-grained Web services in healthcare domain and also for a *Service*

Message Ontology to annotate finer granularity services retrieving meaningful EHR components. These issues are detailed in the following sections.

2.2 Web Service Functionality Ontology

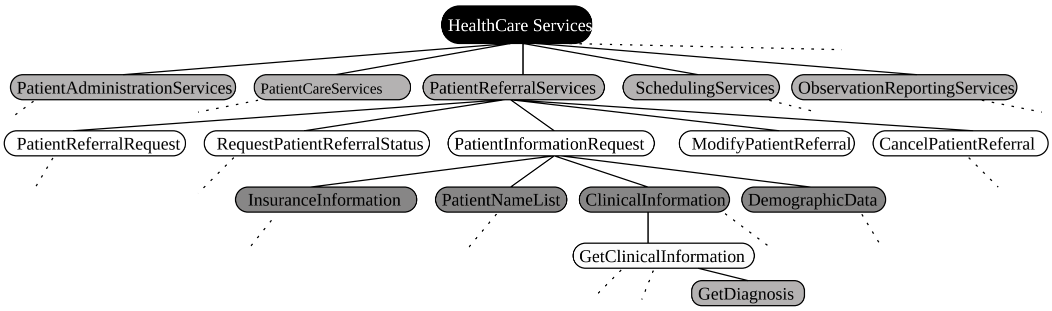

Fig. 2. A Service Functionality Ontology based on HL7

Since HL7 has already been through an effort of categorizing the events in healthcare domain considering service functionality, we propose to use this classification as a basis for a service functionality ontology.

The HL7 standard [23] groups the HL7 events into the following clusters: Patient Administration, Order Entry, Query, Financial Management, Observation Reporting, Master Files, Medical Records/Information Management, Scheduling, Patient Referral, and Patient Care. These clusters also have sub clusters. A partial Web service Functionality Ontology is given in Figure 2 based on HL7 events.

When searching for the right Web services, consumers can consult this ontology to find out services they are looking for by using the functionality semantics of the service. Additionally, service discovery is facilitated by incorporating the nodes of this ontology to the service registry. How this is achieved in UDDI and ebXML registries, is explained in Section 2.7.

It should be noted that our aim is not to propose an ontology but to show how such ontologies, once developed, can be used in semantic mediation. Semantic mediation allows information sources to export their local ontologies to the mediator. If the individual sources have different local ontologies, the mediator needs to reconcile the differences. To enable this reconciliation, the mediator contains mapping rules that explicitly specify the relationships among the ontologies of the different sources. Furthermore, ontology languages contain constructs to state axioms that make it possible to assert subsumption or equivalence with respect to classes and properties in an ontology. These axioms, through reasoners, help discovering further relationships to be used in mapping ontologies.

2.3 Web Service Message Ontology

A Web service in the healthcare domain usually accesses or updates a part of an electronic healthcare record, that is, parts of the EHR constitute the service parameters. An electronic healthcare record may get very complex with data coming from diverse systems such as lab tests, diagnosis, or prescription of drugs which may be in different formats.

As an example, consider the Web service given in Figure 7 Part (b). Although the semantic of action, the "Klinik_Bilgi_Saglayici" service is providing, is clear from the functionality ontology (i.e., it is retrieving clinical information about a patient), it is not clear what the content and format of service parameters like "PatientID" and "ClinicalInformation" are. To achieve interoperability, this additional message semantics is essential and we exploit the EHR based standards in this respect.

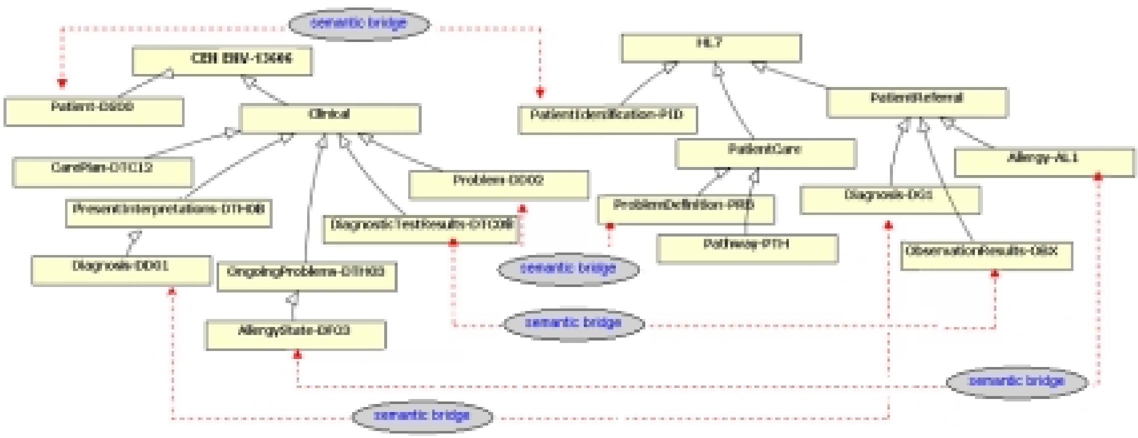

Fig. 3. CEN ENV-13606 and HL7 Clinical Concept Ontologies

Electronic healthcare record (EHR) based standards like HL7 CDA (Clinical Document Architecture) [15], GOM (GEHR Object Model) [3] and CEN's ENV 13606 [7] aim to facilitate the interoperability between Medical Information Systems. However, they do not aim direct machine-to-machine interoperability. Therefore these standards do not prescribe a monolithic EHR architecture; rather they provide conceptual "building blocks" or "meaningful components" by which any clinical model can be represented within the standardized framework. This provides flexibility by allowing the same "building block" to be composed differently by two different institutes, which in turn results in different message structures. This necessitates structural and semantic mappings between the message components in order to automate their interoperation. It is possible to define "clinical concept" ontologies based on the "building blocks" of the EHRs, with ontology definition languages, such as OWL [46]. As an example, in Figure 3, parts of two clinical concept ontologies are presented based on the "building blocks" of HL7 and ENV-13606.

In the Artemis architecture, medical institutions provide Web Services for accessing the components of EHR with a granularity to retrieve the nodes (or the composition of the nodes) of the Clinical Concept Ontologies. The semantics of the service parameters are defined using "message ontologies" which are constructed by using these "clinical concept" ontologies. Once semantically marked up, these elementary Web services are classified under the *Service Functionality* ontology. For example, a Web service retrieving "Diagnosis" information can be classified under "GetClinicalInformation" node as shown in Figure 2.

The medical institutes can develop their own message ontologies to annotate their Web services. However if these ontologies are derived from the Clinical Concept Ontologies based on prominent healthcare standards like HL7, CEN TC251, ISO TC215 and GEHR, then the ontology mapping is facilitated.

2.4 Ontology Mapping

Although representation of the clinical concepts defined by different standardization efforts may result in disparate clinical ontologies initially; defining them through ontology languages opens up the way to mapping them one another through the mapping rules and reasoning.

Consider the two partial clinical concept ontologies from HL7 and ENV-13606 presented in Figure 3. Once such clinical ontologies are defined, the mappings between them can be achieved using the available "Ontology Mappers" such as "MAFRA" [32]. MAFRA uses a component that defines the relations and transformations between ontologies. Generally speaking, ontology mapping has three main dimensions: discovery, representation and execution. Discovery, which is the extraction of the semantic similarity relations between entities of the ontologies, is accomplished by using existing similarity measuring approaches, such as linguistic based algorithms [41]. In Artemis, ontologies are based on well-defined medical informatics standards which facilitate the discovery phase to a great extent.

For representing the similarities in a formal way, MAFRA provides a meta-ontology called Semantic Bridge Ontology (SBO). Semantic Bridges in SBO encapsulate the required information to translate one source entity (concept, relation, property) to a target entity. Semantic Bridges provide mapping cardinalities from 1:1 to m:n, and allow complex structural mappings such as specialization, abstraction, composition and alternatives.

SBO also has concepts to specify conditions, transformation rules, and transformation functions (services) to be used during execution step. It is possible to specify conditions that need to be verified to execute the semantic bridges. Services are used to reference the resources that will be used to handle transformations (i.e. copy an attribute, split a string). SBO is represented in DAML-OIL in MAFRA.

MAFRA has two primitive semantic bridges: Concept Bridge, and Property Bridge. A Concept Bridge defines the semantic equivalence between two ontology classes. At execution step, an instance concept of the target ontology is created for each source concept when the two concepts are related via a concept bridge. In the same way a Property Bridge defines the equivalence between source and target properties.

Once the relationships between two ontologies are defined through “semantic bridges”, the instances of source ontology can be transformed into target ontology instances by evaluating the “semantic bridges” at the execution step [32]. At this step, firstly, the instances of the target ontology are created if the conditions of the related concept bridges evaluate to true. After all instances are created, property bridges are executed and the properties of target instances are set according to them. In the Artemis project, this step is used for converting one healthcare institute's ontology (say, based on ENV-13606) into another (say, based on HL7) by obtaining the necessary "semantic bridges" from the mapping of original ENV-13606 and HL7 based ontologies.

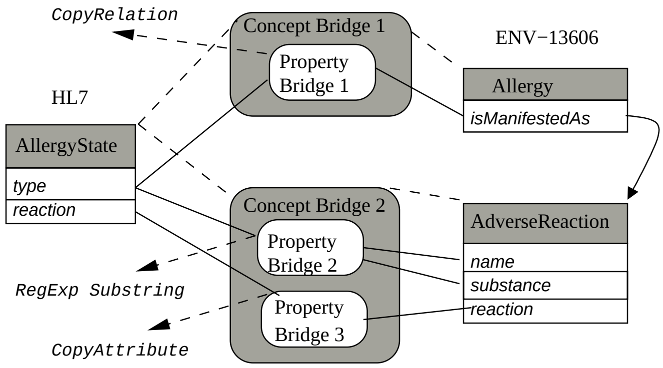

Fig. 4. An Example Mapping Using MAFRA Constructs

As an example, in Figure 4, a mapping using MAFRA constructs is illustrated. In this figure, the "Allergy State" concept of HL7 is mapped to the "Allergy" concept of ENV-13606 through semantic bridges. While a single class is used to represent the "Allergy State" in HL7, the same information is represented with two associated classes, namely "Allergy" and "Adverse Reaction" in ENV-13606. Hence to map these concepts, two "Concept Bridges" are constructed. The "type" attribute in "Allergy State" contains information about "name" and "substance" attributes of "Adverse Reaction". To represent this relation, the "Property Bridge 2" is added to the "Concept Bridge 2". In the execution step this mapping is handled with the help of "RegExp Substring" which is a predefined service of MAFRA, which basically searches/splits a string via regular expressions. The "reaction" attributes in both ontologies which carry the same semantics, are directly mapped through the "Property Bridge 3".

Finally, to express the semantic relation between the "Allergy" and "Adverse Reaction" "Property Bridge 1" is added to the "Concept Bridge 1".

2.5 Exploiting Semantically Annotated Web Services

In this section, we describe how to use the defined semantics to discover the needed Web services.

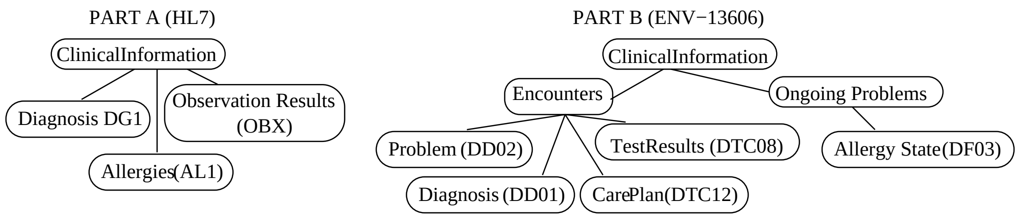

Fig. 5. Clinical Information Representation in two different Systems

Consider Healthcare Institute A, which needs the Clinical Information of a patient stored in Healthcare Institute B. As previously stated Artemis gives the flexibility to the healthcare institutes to define their own clinical message ontologies based on existing standards. Therefore, Healthcare Institute A may define "Clinical Information" as presented in Figure 5 Part A, in terms of the Clinical Concepts defined by HL7 (Figure 3), and Healthcare Institute B may define the same concept, as depicted in Figure 5 Part B, in terms of the Clinical Concepts defined by ENV-13606. In fact these are parts of the "message ontologies" of these institutes, which are used in exchanging "Clinical Information". Notice that both the "building blocks" of these "message ontologies", and also their hierarchical structures are different. Therefore when Healthcare Institute A requests "Clinical Information" of a patient from Healthcare Institute B, both structural and semantic transformation of the documents exchanged are necessary.

The semantic mappings between the concepts in these two "message ontologies" are handled by using Ontology Mappers such as MAFRA to process the "semantic bridges" defined between the "Clinical Concept Ontologies" (i.e. Building blocks of these message ontologies). For example, as shown in Figure 3, "Allergy State (DF03)" in ENV-13606 corresponds to "Allergies (AL1)" in HL7.

If Healthcare Institute B is providing the Web Services for accessing the "Test Results", "Allergy State" and "Diagnosis", after the semantic correspondences between these clinical concepts are determined as explained above, the structural mappings are easily handled as follows: after discovering the Web Services retrieving "Test Results", "Allergy State" and "Diagnosis", the result requested by Healthcare Institute A can be obtained. Here we are assuming that the semantic mappings between these concepts have already been defined through semantic bridges in MAFRA, that is, "Test Results" corresponds to "Observation Results"; "Allergy State" to "Allergies"; and "Diagnosis" to the "Diagnosis" concept. If there are no semantic correspondences between the concepts, the same decomposition process should be applied until finer granularity semantically agreed components are reached.

Since the Web services are annotated with message ontologies based on Clinical Concept Ontologies, it is possible to identify the Web services providing the requested information such as "Diagnosis" from service registries. For this purpose the tModel keys associated with the nodes of Clinical Concept Ontologies are used to find related services in UDDI. In ebXML, a *Clinical Concept* ontology exists (just like the *Service Functionality* ontology) and the related nodes of this ontology such as “Diagnosis:DD01” are used to find the requested services. The details of how this is achieved are presented in Section 2.7. In this way, the information requested in the ClinicalInformation record can be obtained from the Healthcare Institute B as requested by the Healthcare Institute A.

2.6 Semantic Aggregation of Medical Web Services

Although classifying the Web Services through the “semantic category” of the data they are retrieving facilitates the discovery of the services that give a specific part of the EHR data, it may not always be possible to find a service delivering exactly the data requested. For example, a healthcare institute may be requesting “Diagnosis” information whereas the target institute may be providing the diagnosis information as a part of another clinical concept. This may necessitate more complex aggregations of Web Services (such as union, intersection). In other words when we try to compose a Web Service from fine granularity Web services according to the structure and the semantics of the composite Web service output parameter(s), we may not always find disjoint Web services to produce the requested output. Our aim in these sections is to device a methodology for this purpose, that is, gathering an output that can not be delivered by existing Web Services.

As an example consider the case where Healthcare Institute A is requesting Clinical information as shown in Figure 5 Part A but Healthcare Institute B provides Web Services only to retrieve “Encounter” and “Ongoing Problem” information of a patient (Figure 5). Given the semantic structure of “Encounters” and “Ongoing Problems” concepts, it is possible to construct the

"Clinical Information" concept requested by Healthcare Institute A, through a set of *Semantic Aggregation Operators* (SAO). For example we can construct the "Clinical Information" concept of Healthcare Institute A (i.e. ClinicalInformation:A) as follows:
ClinicalInformation:A = (ClinicalInformation:A $\cap_s$ Encounters:B) $\cup_s$ (ClinicalInformation:A $\cap_s$ OngoingProblems:B).

We call the Web Services constructed as "semantic aggregations" of other Web Services as *Virtual Web Services* (VWS). In other words, these virtual Web services are abstractions; they are neither instantiatable nor executable. Rather, they specify how to obtain the required output of a complex Web service from other Web services through *Semantic Aggregation Operators* (SAO).

In other words we are trying to "decompose" a "virtual web service" in terms of existing web services. This is different from the "service composition" models in the literature such as BPEL4WS [5]. The "virtual web services" themselves are not executable directly. In Artemis architecture, the definitions of "virtual web services" are used in obtaining the desired output by using the existing web services through the aggregation operations we define. Through "virtual web services", the system locates the web services necessary to generate the requested output, executes them one by one and constructs the expected result using the *Semantic Aggregation Operators* (SAO).

In the example presented, the Healthcare Institute A uses a *Virtual Web Service* to retrieve the Clinical Information from the Healthcare Institute B.

2.6.1 Semantic Aggregation Operators

We propose a number of “Semantic Aggregation Operations” (SAO) in order to construct *Virtual Web Services* (VWS). These SAOs are as follows:

- $VWS1(\cup_s)VWS2$ *Semantic Union*: This operation is used to construct a VWS which provides the union of the outputs of services VWS1 and VWS2. In other words the output of VWS includes the disjoint concepts provided by VWS1 and VWS2 and the semantically equivalent concepts provided by both only once.

  Example: *Semantic Union* of the VWS whose output semantic is given in Figure 5 Part A with the VWS whose output semantic is given in the same figure Part B produces a VWS whose output semantic is the same as the ontology in Figure 5 Part B. The input semantics of the resultant VWS is defined as the *Semantic Union* of the inputs of the involved Web services.
- $VWS1(\oplus_s)VWS2$ *Semantic Heaping*: This operation can be used to construct a VWS which provides the combined outputs of services VWS1 and VWS2 by collecting all the concepts that take place in both, and disregarding whether the concepts are semantically equivalent or not.

  Example: If we apply *Semantic Heaping* to two VWSs whose output semantics are given in Part A and Part B of Figure 5, the output semantics of the resultant VWS is as given in Figure 6. The input semantics of the resultant VWS is defined as the *Semantic Union* of the inputs of the involved Web services.

- $VWS1(\cap_s)VWS2$ *Semantic Intersection*: This operation can be used to construct a VWS which provides the semantically equivalent concepts provided by both VWS1 and VWS2.

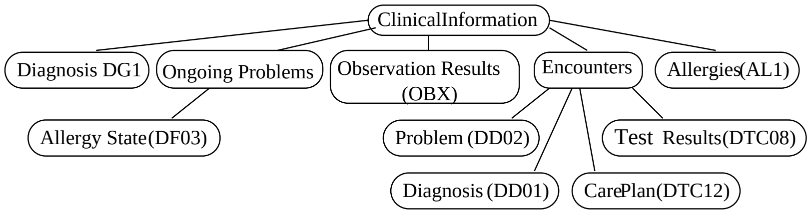

Fig. 6. Semantic Heaping Example

- $VWS1(\ominus_s)VWS2$ *Semantic Difference*: This operation can be used to construct a VWS which gives the concepts provided by VWS1 excluding the semantically equivalent concepts provided by VWS2.

  Example: Semantic Difference of a VWS whose output semantic is given in Figure 5 Part A from a VWS whose output semantic corresponds to the "Ongoing Problems" node of the ontology presented in the same figure Part B will result a VWS producing the only "Diagnosis (DG1)" and "Observation Results (OBX)" nodes of the ontology depicted in Part A.

- $VWS1(\gg_s)VWS2$ *Semantic Contain*: This operation can be used to check whether the concepts provided by VWS1 is a superset of the concepts provided by VWS2.

  Example: If we use the *Semantic Contain* operation between the VWS whose output semantic is given in Figure 5 Part B and the VWS whose output semantic is given in the same figure Part A, we get a "true" answer since all of the concepts produced by the first VWS are also produced by the second one.

Given these "semantic aggregation" operators, coarse grained Web Services can be composed from finer granularity services even when there is no finer granularity service retrieving exactly the requested data. For example, as shown in Figure 5, the Web Service providing "Encounters" Information of a patient is classified with the "clinical concepts" in its output such as Problem:DD02, TestResults:DTC08, Diagnoses:DD01, and CarePlan:DTC12 (Figure 5). Hence this Web Service is a candidate for aggregation in order to gather the data requested by Healthcare Institute A.

The results of these aggregations, i.e. Virtual Web Services, are also re-usable components. They are inserted as instances into the *Service Functionality* ontology together with their descriptions. For instance, the virtual service retrieving Clinical Information of a patient in the running example is stored as an instance of the "GetClinicalInformation" node of the *Service Functionality* ontology presented in Figure 2. Whenever these two hospitals interact again, these VWS definitions can be reused.

Providing such Virtual Web Services and creating a repository from these VWS, in the long run, may improve the interoperability of Medical Information Systems. Although these VWS may seem as bilateral agreements between institutes, they can be used to create a wider community through transitive agreements as discussed in [1].

2.7 Relating Web Service Ontologies with Web Service Registries

Once the semantics of Web services are specified, it is necessary to relate these with the services advertised in service registries.

There are two key issues in this process: the first one is where to store the ontologies. UDDI does not provide a mechanism to store an ontology internal to the registry. ebXML, on the other hand, through its classification hierarchy mechanism allows domain specific ontologies to be stored in the registries. Note that for UDDI registries, domain specific ontologies can be stored by the standard bodies who define them and the server, where the service is defined, can host the semantic description of the service instance.

The second key issue is how to relate the services advertised in the registry with the semantics defined through an ontology. The mechanism to relate semantics with services advertised in the UDDI registries is to use the tModel keys and the category bags of registry entries. tModels provide the ability to describe compliance with taxonomies, ontologies or controlled vocabularies. Therefore if tModel keys are assigned to the nodes of the ontology (for example given in Figure 2) and if the services put the corresponding tModel keys in their category bags, it is possible to locate services conforming to the semantics given in a particular node of this ontology. This issue is elaborated in [10].

An ebXML registry [16], on the other hand, allows to define semantics basically through two mechanisms: first, it allows properties of registry objects to be defined through "slots" and, secondly, metadata can be stored in the registry through a "ClassificationScheme". Furthermore, "Classification" objects explicitly link the services advertised with the nodes of a "ClassificationScheme". This information can then be used to discover the services by exploiting the ebXML query mechanisms.

Consider for example the service Functionality Ontology given in Figure 2. Such a hierarchy can be stored in an ebXML registry through the piece of code as shown in Figure 7 Part (a), and then the registry objects can be related with the nodes in the hierarchy. In this way it is possible to give meaning to the services. In other words, by relating a service with a node in the classification hierarchy, we make the service an explicit member of this node and the service inherits the well-defined meaning associated with this node as well as the generic properties defined for this node. As an example, assume that there is a service instance in the ebXML registry, namely, "Klinik_Bilgi_Saglayici".

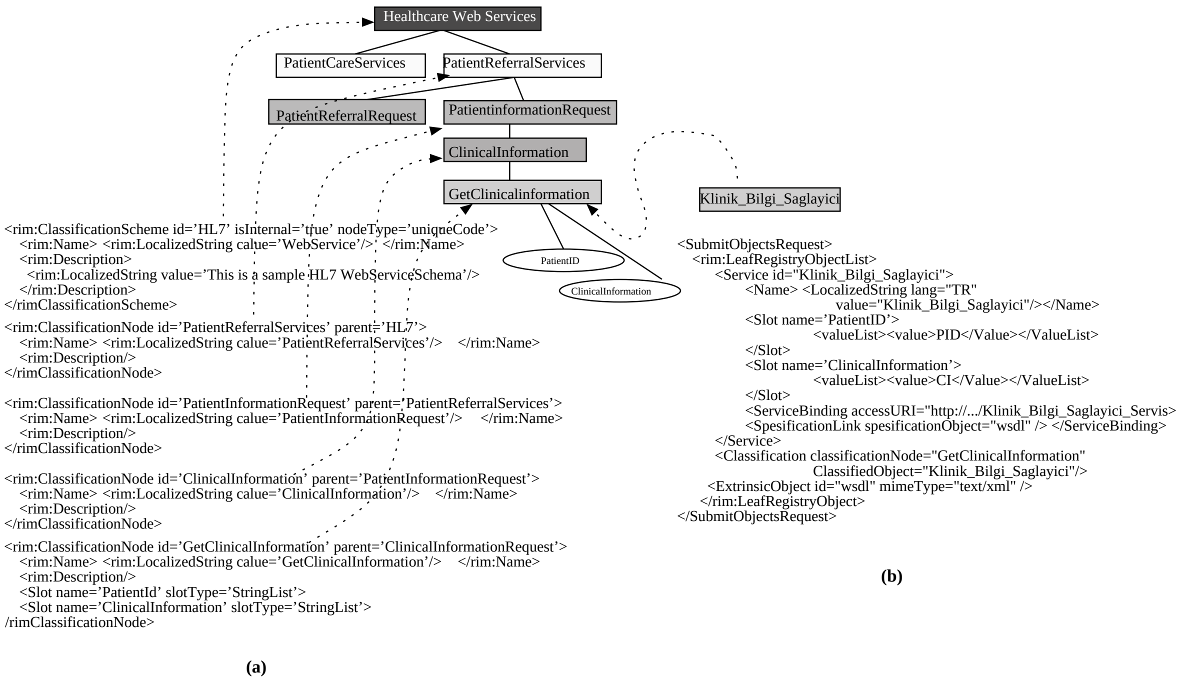

Fig. 7. Defining Ontology Classes in ebXML and Relating a Service Instance with the Ontology Class

When we associate "Klinik_Bilgi_Saglayici" with the "GetPatientClinicalInformation" node through a "SubmitObjectsRequest" as shown in Figure 7 Part (b), its meaning becomes clear: this service is providing patient clinical information. Furthermore, the "Klinik_Bilgi_Saglayici" service inherits properties of the "GetPatientClinicalInformation" service such as "PatientID" and "ClinicalInformation".

Finally, how to store OWL ontologies into ebXML registries and how to associate these ontologies with Web services are described in [13], [14].

3 System Architecture

The Artemis project addresses the interoperability problem in the healthcare domain where organisations have proprietary application systems to access data. To exchange information there are different standards like HL7, GEHR or CEN's ENV 13606. The aim of the Artemis project is to allow organisations keep their proprietary systems, yet expose the functionality through Web services. Furthermore, we propose an ontology based description of these data exchange standards. One of the goals of using ontologies is to reduce (or to eliminate) conceptual and terminological differences among the healthcare data exchange standards through semantic mediation.

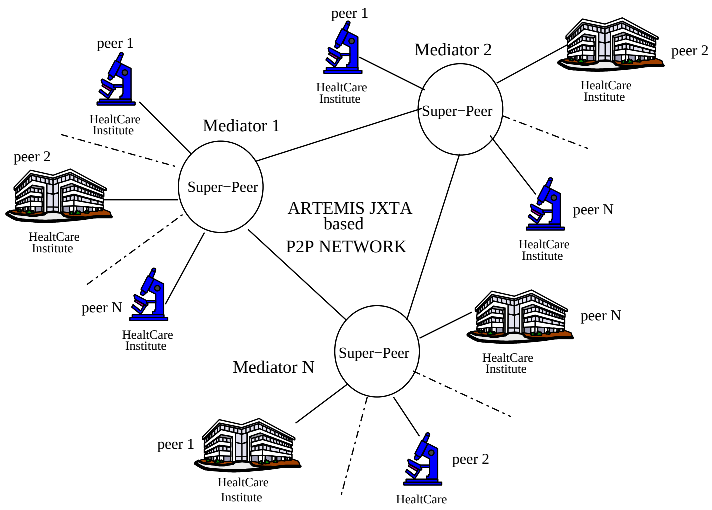

Fig. 8. Artemis P2P Architecture

Mediators are developed to process data from possibly several data sources and to prepare them for the effective use by applications [49]. However with WWW becoming the global communication medium and with the Semantic Web initiative, ontologies are becoming the primary part of the mediation process.

Artemis Web service architecture does not rely on globally agreed ontologies: rather healthcare institutes develop their own ontologies. However, it is reasonable to expect healthcare institutes to develop their own ontologies or use already defined ones based on the concepts provided by the existing healthcare information standards since considerable semantic information is already captured there.

Artemis architecture then helps to reconcile the semantic differences among healthcare institutes through the mediator component. To provide scalability and discovery of other mediators, it has a P2P communication architecture. An overview of the Artemis architecture is given in Figure 8.

3.1 Artemis Mediator P2P Architecture

In Artemis, healthcare institutes communicate with each other through mediators which resolve their differences bilaterally. When it comes to how to organize the mediators we make the following observations:

- The mediators must have a distributed architecture to provide for scalability.
- When a healthcare institute, say A, wants to communicate with another healthcare institute, say B, it should be possible to automatically locate the mediator of B.
- There are efficiencies to be gained by logically grouping the healthcare institutes which communicate often through a single mediator.

With these considerations in mind, Artemis mediators are designed and implemented as JXTA super peers. JXTA is an Open Source project [31] supported and managed by Sun Microsystems. Basically, JXTA is a set of XML based protocols to implement typical P2P functionalities. In the JXTA super peer based architecture, peers in a peer group communicate with their super peer to advertise their capabilities as well as to search for other peers.

In Artemis, each mediator is a super peer serving the healthcare institutes in its logical peer group. Super-peers employ semantics based routing indices where semantics are used to locate the healthcare institutes. On registration, the peer provides this information to its super-peer.

3.2 Artemis Mediator Component

Generally speaking, semantic mapping is the process where two ontologies are semantically related at conceptual level and source ontology instances are transformed into target ontology entities according to those semantic relations. In Artemis, the source and target ontologies belong to the two healthcare institutes willing to exchange information. However, the mapping of these two ontologies is achieved through the reference ontologies stored in the mediator: the generic *Service Functionality* and *Service Message* ontologies. The mediator resolves the semantic differences between source and target ontologies by using these and the Semantic Bridge ontologies.

It should be noted that since all the ontologies involved are somehow related with the basic healthcare standards, the mediation process is simpler and hence more efficient. Furthermore, resolved semantic differences are stored as *Virtual Web Services (VWS)* to be reused as explained in Section 2.6.

The mediator architecture, which is shown in Figure 9, has the following subcomponents:

- Ontology server: The Ontology server contains the following ontologies:
  - *Service Functionality* and *Service Message* ontologies: Each healthcare institute may develop its own Service Functionality and Service Message

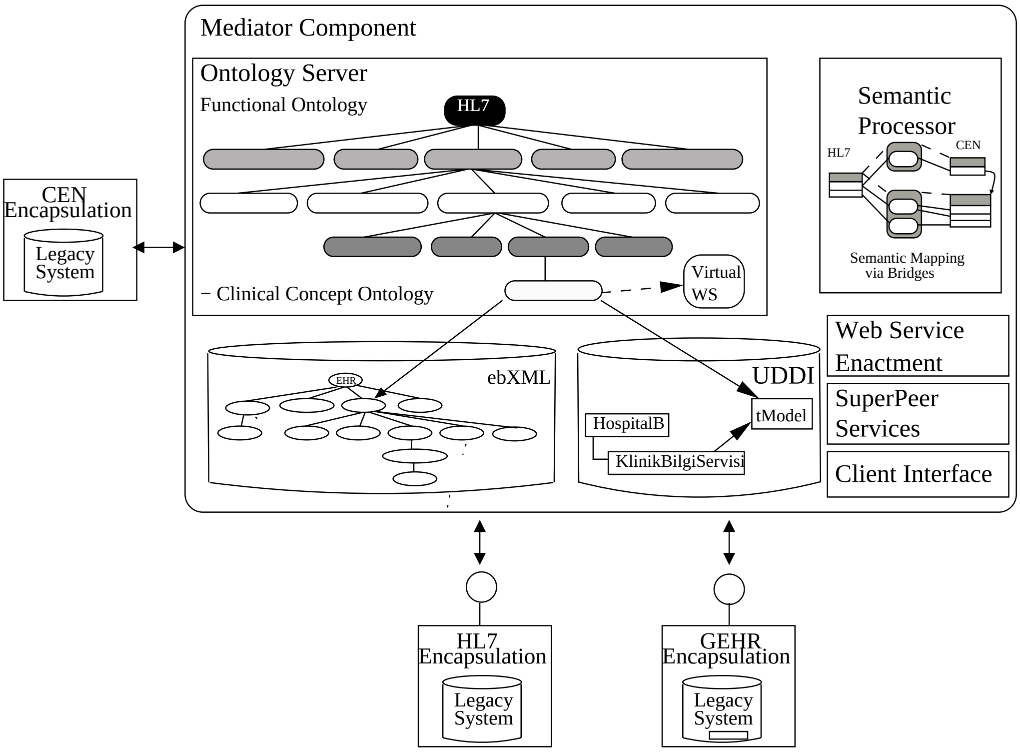

Fig. 9. An Overview of the Mediator

ontologies based on existing healthcare information standards. The minimum requirement is annotating their services through such ontologies.

· *Virtual Web Services* subsystem handles the creation of Virtual Web Services (VWSs) to provide complex aggregations of Web services. The creation of VWSs is realised according to the mappings between the ontologies of Web services' input and output semantics. Newly created VWSs are classified according to the *Service Functionality Ontology* of the requesting party for its possible future reuse.

- *Semantic Processor*: There may be more than one *Service Functionality* and *Service Message* ontologies in the mediator and the mediator generates the mappings between them using its own reference ontologies based on the healthcare standards. In Artemis, MAFRA is used to represent the mappings and to transform the ontology instances. MAFRA uses the Semantic Bridge Ontology to define the mappings and includes a transformation engine. The mediator stores the previously defined mappings via semantic bridges. For example, the semantic equality relation between the "DiagnosticTestResult" concept in ENV 13606, and the "ObservationResult" concept in HL7 can be represented using MAFRA semantic bridges as follows:

```
<a:ConceptBridge rdf:ID="CB163312">
 <a:relatesTargetEntity rdf:resource=
  "http://www.srdc.metu.edu.tr/HL7#ObservationResult"/>
 <a:relatesSourceEntity rdf:resource=
  "http://www.srdc.metu.edu.tr/CEN#DiagnosticTestResult"/>
 <a:abstract rdf:resource="&a;True"/>
</a:ConceptBridge>
```

Note that, more complex mappings can be represented using "semantic bridges", such as compositions, alternatives, and transformations aided by external functions.

At runtime the source ontology instances are transformed into target ontology instances by providing the source instance and the RDF [33] representation of mapping to the transformation engine of MAFRA.

- Service registries like UDDI and ebXML: The Web services of the involved healthcare institutes are published in the UDDI or ebXML registries of the mediator.
- Web service Enactment Component handles the invocation of the Web Services and transmits the results of the Web Services.
- Superpeer Services Component contains the services that provide the communication with other Mediators in a P2P infrastructure. Basically, these services implement the JXTA Protocols. For example, Discovery Service that implements the JXTA Peer Discovery Protocol is used to find other Mediators through a semantics based search mechanism.
- Client Interface handles the communication of healthcare institutes with

the mediator using client-mediator protocol.

3.3 Implementation Status of the Proposed Architecture

A proof of concept prototype of the proposed system is developed within the scope of the Artemis project. This prototype is implemented in Java 2 [28]. As the P2P Mediator architecture, JXTA 2 [31] is used. A number of JXTA Protocols have been implemented for providing the communication between the Mediators. For example, a Discovery Service is implemented over the JXTA Peer Discovery Protocol to find the other Mediators through a semantics based search mechanism.

The "Functionality, Message and Clinical Concept Ontologies" have been developed using Protege with OWL plugin [40]. In the "Semantic Processor Component", JENA 2.0 [29] is used for parsing the ontologies. MAFRA [32] is used for defining the mappings between ontologies, moreover MAFRA engine has been adapted in order to call it from a Java class.

The Medical Web Services are published to UDDI and ebXML registries using the UDDI 2.0 conformant Java Web Services Developer Pack [30] and ebXML RIM v2.1 conformant OASIS ebXML Registry Reference Implementation Project (ebxmlrr) [17] respectively. Finally in the "Web Service Enactment Component", the JXTA-SOAP Bridge [4] is used to deploy and invoke Web services in the JXTA environment.

4 Related Work

Currently, describing the semantics of Web services is a very active research area. DAML-S [8] (later OWL-S) is a comprehensive effort defining an upper ontology for Web services. Service discovery through DAML-based languages is also addressed in the literature [9,34,35,38] where artificial intelligence techniques are used to discover services.

In [37], an RDF mapping meta-ontology, called RDF Translation (RDFT), is proposed which specifies a language for mapping XML DTDs to and from RDF Schemas for business integration tasks.

In ChattyWeb [1], the P2P paradigm is used to improve semantic interoperability, in particular in revealing new possibilities on how semantic agreements can be achieved. It is argued that establishing local agreements is a less challenging task than establishing global agreements by means of globally agreed schemas or shared ontologies. Once such local agreements exist, through the "semantic gossiping" process proposed, global agreements can be achieved in a P2P manner.

The work described in this paper has benefited from the previous work in the following areas:

- *Semantic Web Service Architecture:* The semantic architecture of Artemis is adapted from the Semantic Web Enabled Web Services (SWWS) [6] architecture. The SWWS architecture considers semantics as a vertical layer that may be exploited by the horizontal layers of the Web service stack such as service description (including the documents exchanged), publishing, discovery as well as service flow and composition as shown in Figure 10 [6]. The authors describe how semantics can be exploited in different levels of the Web service stack and stress the importance of ontologies and semantic mediation to deal with the interoperability problem.

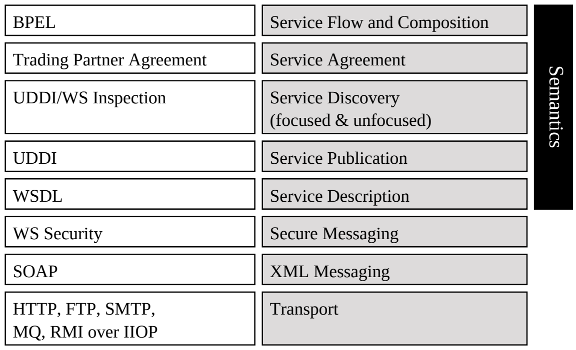

Fig. 10. Web Service Stack and Semantic

A detailed overview of the Web Service Modeling Framework (WSMF) is given in [18].

- *Ontology Mapping*: The ontology mapping component of Artemis mediator uses the technologies described in [20] where a semantic mapping and reconciliation engine is developed within the scope of the Harmonise project [19]. The Harmonise project aims to develop a harmonization network for the tourism industry to allow participating tourism organisations to keep their proprietary data format and use ontology mediation while exchanging information in a seamless manner. For this purpose they have defined a *Interoperability Minimum Harmonization Ontology* and an interchange format for tourism industry. The MAFRA [32] tool is used for ontology mediation.
- Extending the UDDI registries with semantic capabilities is addressed in [10,11], where we describe a mechanism to relate DAML-S ontologies with services advertised in the UDDI registries. [39] also addresses importing semantic to UDDI registries where DAML-S specific attributes such as *inputs, outputs* and *geographicRadius* are represented using tModel mechanisms of

UDDI.

- A semantic-based Web service composition facility for ebXML Registries is described in [12]. [13] describes how ebXML registries can be enriched with OWL ontologies for efficient Web service discovery.

5 Conclusions

We believe that Web service semantics need to be addressed in a domain specific way since different domains have very different needs; they have evolved very differently and the semantics is domain specific information. In this paper we present a novel architecture for exploiting the Web service semantics in the healthcare domain. The work described is being realized within the scope of Artemis project (IST-2103) [2].

The contributions of this work are as follows:

- To provide interoperability in the healthcare domain, we expose healthcare applications by wrapping them as semantically enriched Web services. To the best of our knowledge, Artemis is the first initiative to use semantically enriched Web services in the healthcare domain. In fact, only very recently Web services started to appear in the medical domain. An important industry initiative to use Web services is "Integrating the Healthcare Enterprise (IHE)" [26]. IHE has defined a few basic Web services such as "Retrieve Information for Display Integration Profile (RID)". Yet, since IHE does not address semantic issues: to use IHE Web services, it is necessary to conform to their specification exactly, by calling the Web services with the names they have specified and providing the messages as instructed in its specification.

- Another contribution of the work described in this paper is identifying the need for service functionality and service message ontologies to semantically annotate Web services. The semantic information required is based on the existing healthcare standards. HL7 constitutes the basis of *Service Functionality Ontology* since HL7 has categorized the events in healthcare domain by considering service functionality.

  We organize the "meaningful components" defined by the electronic healthcare record (EHR) based standards like HL7 CDA (Clinical Document Architecture) [15], GOM (GEHR Object Model) [3] and CEN TC251's ENV 13606 [7] into ontologies. We later use such clinical ontologies in associating semantics with the messages and documents exchanged through the Web services.

- Although we propose ontologies based on the prominent healthcare standards, the ontologies we are proposing are just to facilitate ontology mediation. In the Artemis architecture, the healthcare institutes can develop their own ontologies. However, when these ontologies are based on standards developed by the healthcare standardization bodies like CEN TC251, ISO TC215, GEHR or HL7, we show that ontology mappings are facilitated to a great extend through semantic mediation.

  The mediator architecture in Artemis is based on a peer-to-peer infrastructure to provide scalability and to facilitate the discovery of other mediators.

- Although classifying the Web Services through the "semantic category" of the data they are providing facilitates the discovery of the services fetching a specific part of the EHR data, it may not always be possible to find a service delivering exactly the data requested. For example, a healthcare institute may be providing the diagnosis information as a part of another clinical concept. This may necessitate more complex aggregations of Web Services. We address how complex aggregation of Web services can be handled by taking advantage of *Semantic Aggregation Operators* that we have defined.

- A proof of concept prototype of the system is developed to guide the full scale implementation which is underway.

In this paper, we mainly focused on the clinical concept part of the message ontologies. Our main motivation for concentrating on clinical concept ontologies is that the electronic healthcare record based standards present detailed semantics in this regard. However healthcare is a many-to-many business. It is not only connecting a hospital to its branch clinics but to an array of internal and external agencies such as insurance entities, financial institutes and government agencies. Therefore there are other aspects of healthcare informatics such as billing and insurance that need to be covered. Our future work includes extending message ontologies with semantic concepts to handle these aspects including financial information.

References

[1] K. Aberer, P. Cudre-Mauroux, M. Hauswirth, "The Chatty Web: Emergent Semantics Through Gossiping", The Twelfth International World Wide Web Conference, WWW2003, Budapest, Hungary, May 2003

[2] Artemis Project, http://www.srdc.metu.edu.tr/webpage/projects/artemis

[3] T. Beale, GEHR Object Model Architecture (GOM), http://www.gehr.org/technical/model_architecture/model_architecture_4.1E.pdf

[4] Kevin A. Burton, Bridge, http://soap.jxta.org/servlets/ProjectHome

[5] Business Process Execution Language for Web Services (BPEL4WS), http://www-106.ibm.com/developerworks/library/ws-bpel/

[6] C. Bussler, D. Fensel, A. Maedche, “A Conceptual Architecture for Semantic Web Enabled Web Services”, SIGMOD Record, Vol. 31, No. 4, December 2002.

[7] CEN TC/251 (European Standardization of Health Informatics) ENV 13606, Electronic Health Record Communication, http://www.centc251.org/

[8] DAML-S, DAML Services Coalition, in Proceedings of the International Semantic Web Working Symposium (SWWS), July 2001.

[9] G. Denker, J. R. Hobbs, S. Narayan, R. Waldinger, “Accessing Information and Services on DAML-Enabled Web”, Semantic Web Workshop, Hong Kong, China, 2001.

[10] A. Dogac, G. B. Laleci, Y. Kabak, I. Cingil, “Exploiting Web Service Semantics: Taxonomies vs. Ontologies”, IEEE Data Engineering Bulletin, Vol. 25, No. 4, December 2002.

[11] A. Dogac, I. Cingil, G. B. Laleci, Y. Kabak, “Improving the Functionality of UDDI Registries through Web Service Semantics”, 3rd VLDB Workshop on Technologies for E-Services (TES-02), Hong Kong, China, August 23-24, 2002.

[12] A. Dogac, Y. Kabak, G. Laleci, ”A Semantic-Based Web Service Composition Facility for ebXML Registries”, 9th International Conference of Concurrent Enterprising, Espoo, Finland, June 2003.

[13] A. Dogac, Y. Kabak, G. Laleci, “Enriching ebXML Registries with OWL Ontologies for Efficient Service Discovery”, in Proc. of RIDE’04, Boston, March 2004.

[14] A. Dogac, Y. Kabak, G. Laleci, C. Mattocks, F. Najmi, J. Pollock, “Enhancing ebXML Registries to Make them OWL Aware”, Submitted to the Distributed and Parallel Databases Journal, Kluwer Academic Publishers. http://www.srdc.metu.edu.tr/webpage/publications/2004/_DAPD_ebXML-OWL.pdf

[15] R. H. Dolin, L. Alschuler, S. Boyer, and C. Beebe, “An Update on HL7’s XML-based
Document Representation Standards”, http://www.amia.org/pubs/symposia/D200113.PDF

[16] ebXML, http://www.ebxml.org/

[17] OASIS ebXML Registry Reference Implementation Project (ebxmlrr) http://ebxmlrr.sourceforge.net/

[18] D. Fensel, C. Bussler, “The Web Service Modeling Framework WSMF”, Electronic Commerce Research and Applications, Vol. 1, Issue 2, Elsevier Science B.V., Summer 2002.

[19] Harmonise Project, IST-2000-29329, Tourism Harmonisation Network, http://www.harmonise.org/

[20] Harmonise Project, Deliverable 3.2: Semantic mapping and Reconciliation Engine subsystems, http://sourceforge.net/projects/hmafra

[21] The Good Electronic Health Record, http://www.gehr.org

[22] A. Gupta, , B. Ludascher, M. Martone, “Registering Scientific Information Sources for Semantic Mediation”, 21st Intl. Conf. on Conceptual Modeling, (ER2002), Finland, October 2002.

[23] Health Level 7 (HL7), http://www.hl7.org

[24] HL7, Chapter 11 Patient Referral, http://www.hl7.org/library/General/v231.zip

[25] ICD-10, International Statistical Classification of Diseases and Related Health Problems, http://www.who.int/whosis/icd10/

[26] Integrating the Healthcare Enterprise, www.rsna.org/ihe

[27] ISO TC215, International Organization for Standardization, Health Informatics Technical Committee, http://www.iso.ch/iso/en/stdsdevelopment/tc/tclist/TechnicalCommitteeDetailPage.TechnicalCommitteeDetail?COMMID=4720

[28] Java 2 Enterprise Edition 1.4, http://java.sun.com/j2se

[29] JENA, A Semantic Framework for Java, http://jena.sourceforge.net/

[30] Java Web Services Developer Pack, http://java.sun.com/webservices/webservicespack.html

[31] Project JXTA, http://www.jxta.org/

[32] A. Maedche, D. Motik, N. Silva, R. Volz, "MAFRA-A MApping FRAmework for Distributed Ontologies", In Proc. of the 13th European Conf. on Knowledge Engineering and Knowledge Management EKAW-2002, Madrid, Spain, 2002.

[33] Resource Description Framework (RDF), http://www.w3.org/RDF/

[34] S. A. McIlraith, T. C. Son, H. Zeng, "Semantic Web Services", IEEE Intelligent Systems, March/April 2001, pp. 46-53.

[35] S. A. McIlraith, T. C. Son, H. Zeng, "Mobilizing the Semantic Web with DAML-Enaled Web Services", Semantic Web Workshop 2001, Hongkong, China.

[36] E. Motta, J. Domingue, L. Cabral, M. Gaspari, "IRS II: A Framework and Infrastructure for Semantic Web Services", 2nd International Semantic Web Conference, Florida, USA, October 2003.

[37] B. Omelayenko, D. Fensel, C. Bussler, "Mapping Technology for Enterprise Integration", Proc. of Special Track on Semantic Web at The 15th International FLAIRS Conference, USA, May 2002.

[38] M. Paolucci, T. Kawamura, T. Payne, K. Sycara, "Semantic Matching of Web Services Capabilities", in Proc. of Intl. Semantic Web Conference, Sardinia, Italy, June 2002.

[39] M. Paolucci, T. Kawamura, T. Payne, K. Sycara, "Importing the Semantic Web in UDDI", in Web Services, E-Business and Semantic Web Workshop, 2002.

[40] Protege OWL Plugin, http://protege.stanford.edu/plugins/owl/

[41] E. Rahm, P. A. Bernstein, "A survey of approaches to automatic schema matching", VLDB J. Vol. 10, No. 4, Dec. 2001.

[42] N. Silva, J. Rocha, "Semantic Web Complex Ontology Mapping", Proceedings of the IEEE Web Intelligence 2003 conference, Canada, October 2003.

[43] SNOMED Clinical Terms, http://www.snomed.org/snomedct_txt.html

[44] Simple Object Access Protocol (SOAP), http://www.w3.org/TR/SOAP/

[45] UDDI: Universal Description, Discovery and Integration, http://www.uddi.org, 2001.

[46] Web Ontology Language OWL, http://www.w3.org/TR/owl-features/

[47] Web Service Description Language (WSDL), http://www.w3.org/TR/wsdl

[48] WS-Attachment, http://www-106.ibm.com/developerworks/webservices/library/ws-attach.html

[49] G. Wiederhold, "Mediators in the Architecture of Future Information Systems", IEEE Computer, Vol.25 No.3, March 1992.

6 Figures