PART B

COLLABORATIVE PROJECT

Table of Contents

B1. CONCEPT AND OBJECTIVES, PROGRESS BEYOND STATE-OF-THE-ART, S/T METHODOLOGY AND WORK PLAN

B1.1 Concept and project objective(s)

Interoperation between electronic services, and more generally the business processes embodied by services, is one of the most challenging and pressing issues in today's increasingly globalized and de-centralized economy. This proposal puts forward a unified research program that will enable a substantial simplification in the establishment and maintenance of *service collaborations*, that is, families of services, possibly from different organizations, that are interoperating to achieve both individual and shared objectives. Based on an innovative foundation, the research will develop scientific advances, techniques, and tools to dramatically simplify the design and deployment of infrastructure to support service collaborations, the ability of services to join such collaborations, and the evolution of such collaborations as the marketplace and competitive landscape change.

The over-arching objective of the proposed research is to dramatically reduce the effort and time-to-usage of designing, deploying, and maintaining environments that support service collaborations.

The name for this project, *Artifact-Centric Service Interoperation* (*ACSI*) is based on the two fundamental constructs used in our new framework, namely, *interoperation hub* and (*dynamic*) *artifact* (also known as "business artifact"). An **interoperation hub** serves as a virtual rendezvous for multiple services that are to work together towards a common goal. Domain-specific interoperation hubs such as the EasyChair conference submission management system or Salesforce.com have already shown the value of interoperation hubs; the research here will make it possible to easily create, launch, participate in, and maintain ACSI interoperation hubs in essentially any application domain. Similar to EasyChair, an ACSI interoperation hub will serve as the anchor for a *collaboration environment*, that is, an IT environment used to support large numbers of service collaborations that operate independently, but which focus on essentially equivalent common goals. Unlike orchestrators, an interoperation hub works well in the context of open service networks. These hubs are primarily reactive, serving as a kind of *structured white board* that participating services can refer to, that can be updated with information relevant to the group, that can assist the services by carrying out selected tasks, and that can notify services of key events.

The interoperation hubs used here will be structured around **dynamic artifacts**. These provide an holistic marriage of data and process, both treated as first-class citizens, as the basic building block for modeling, specifying, and implementing services and business processes. In the context of single enterprises, IBM has shown that the use of artifacts can lead to substantial cost savings in the design and deployment of business operations and processes, and can dramatically improve communication between stakeholders, especially in cases where they represent different "silos" of the enterprise. Artifacts can give an end-to-end view of how key conceptual business entities evolve as they move through the business operations, in many cases across two or more silo's. As a result, *artifacts can substantially simplify the management of "hand-off" of data and processing between services and organizations*. A key pillar of the ACSI research is to generalize the advantages of artifacts to the context of interoperation hubs and service collaborations. Very importantly, while the interoperation hubs themselves will take advantage of the artifact paradigm, the services participating with such hubs are *not required* to be artifact-centric; they can be conventional SOA services (including legacy services with SOA adapters).

The ACSI research program will develop targeted scientific foundations, techniques, and tools that will simplify the creation and maintenance of collaboration environments in the following measurable ways:

1. **At least 40% reduction, over conventional techniques, in the design and deployment of environments that support large numbers of service collaborations with similar goals**
2. **At least 20% reduction, over conventional techniques, in the costs of on-boarding into, and maintaining, service collaborations**
3. **At least 30% reduction in on-going manual activity needed to support typical service collaborations**
4. **At least 90% of data transformation in service collaborations will be automated**

The ACSI research program is based on three tightly coupled activity streams. The ***scientific stream*** develops a new framework for enabling flexible service interoperation that is highly scalable and supportive of long-tail contexts, and develops a variety of techniques and tools to enable better design for service collaborations, rich discovery of services and service capabilities, and easy on-boarding of services into collaborations. The techniques include extending automatic verification and process mining techniques to incorporate data as found in dynamic artifacts, establishing a logic-based foundation for specifying transformations and evolution of artifact schemas, and auto-generation of adapters and other components for collaboration environments based on interoperation hubs. The ***technology stream*** will combine the results of the scientific stream into a coherent prototype system (the ACSI Hub System), which allows for establishing environments that can support multiple service collaborations, and enables demonstration and testing of the scientific results. Much of the tool development and prototyping will be done in an open-source context. ACSI interoperation hubs will be deployable as Software as a Service (SaaS). This will enable economies of scale when establishing a collaboration environment, and permit a cost model based on usage rather than equipment or software ownership. Finally, the ***validation stream*** will focus on the application, testing, and refinement of the prototype system against two real-world scenarios.

This proposal will address several of the main components of Objective ICT-2009.1.2 (b) – "Highly Innovative Service/Software Engineering". The development of the ACSI interoperation hub framework, based on the end-to-end artifact-centric approach, is a highly innovative approach to supporting service collaborations that can be applied to very large, dynamic open service networks. The use of artifacts in the hub, and techniques for verification, automatic discovery, and auto-generation of adapters helps to support user development of services, and systems evolvability and acquisition. The foundations work on reasoning about behaviours, which is analogous to and will be combined with Description Logics, will enable reasoning and incorporation of domain knowledge in all phases of the service lifecycle. The dynamic artifacts used to define the behaviour of the interoperation hubs are specified in high-level executable languages. The logic-based foundations, the verification research, and the process mining research will all help to lay the basis for, and in some cases within the three year project enable, automatic support at run-time for decisions and changes that are currently adopted at design time. The scientific stream will develop (formal and systems-level) verification and validation methods, tools and techniques assuring the quality of open, large-scale, dynamic service systems without fixed system boundaries, addressing the complete service lifecycle.

The project team includes world-class researchers in all of the key technical areas needed for this research, including experts on artifact-centric business processes, verification, data integration, ontologies for data access, process mining, services architectures, and business process management.

In the following, Subsection 1.1.1 overviews and motivates the notions of interoperation hub and dynamic artifact. Subsection 1.1.2 describes the Scientific Stream; Subsection 1.1.3 describes the Technology Stream; and Subsection 1.1.4 describes the two pilot scenarios of the Validation Stream. And Subsection 1.1.5 summarizes how the ACSI results address many of Objective ICT-2009.1.2 (b) target outcomes.

1.1.1 Interoperation hubs and dynamic artifacts

The first of seven targeted objectives for the ACSI research is:

Objective 1: Development of the ACSI interoperation hub framework, based on the use of dynamic artifacts, and that provides a formal foundation, and can incorporate targeted scientific advances in support of access privileges, verification, process mining, auto-generation, and evolution.

This will be achieved by building on the notions of dynamic artifacts and interoperation hubs, which are described in this subsection. Preceding that, a brief review of conventional technology for supporting service collaborations is given. The subsection concludes by highlighting the ways that the two notions, taken together and extended by the ACSI research, can substantially simplify the design, deployment, and maintenance of service collaborations, enable the use of a high-level yet executable language, and support scalability and long-tails in an open service network.

The context. There is a great need for efficient, effective service collaborations in many sectors, including commerce, energy distribution, government services, supply chains, and healthcare. Whether within a single organization, or working across multiple organizations, the current approaches to designing and deploying business operations and processes to support organizational goals are based on a family of complementary approaches, including Service Oriented Architecture (SOA), Business Process Management (BPM), Master Data Management (MDM), and Business Activity Monitoring (BAM). These approaches use a variety of different conceptual models, each of which decomposes the problem space in different ways, significantly adding to the complexity of establishing and maintaining the solution. A key pillar of the ACSI research is to base service collaborations on a single, unifying conceptual model, namely on dynamic artifacts. Artifacts enable an end-to-end view of a service collaboration that cuts across the participating services, and shows clearly how each service fits into the collaboration, in terms of both data manipulated and tasks performed. This makes it easy to understand how each service is contributing to the overall objectives. Further, most other conceptual models used in BPM naturally fit around the artifact model.

Turning to architectures that support interoperation, the two dominant approaches today are orchestration and choreography. Neither is targeted at enabling scalable service collaborations in open networks and with support for the long tail. For one thing, both are focused primarily on message passing, which tends to obscure how data of common interest evolves through the lifecycle of the end-to-end shared process. Further, they are not well equipped to handle variations in a principled manner. As a result, it is difficult to widely re-use an orchestrator, or the family of constraints characterizing a choreography.

Interoperation hubs. A central pillar of the ACSI research program is to substantially develop a new framework for facilitating collaboration between services, based on the notion of interoperation hub (initial ideas can be found in [HNN09], written by the ACSI project PIs from IBM Research). We shall illustrate this notion, and also the notion of artifact, using an example based on Hiring within a single corporation. (Generalization to supporting multiple organizations shall be discussed shortly.) This example was chosen because it is reasonably familiar to most readers. However, the pilot stream of the proposed research will focus on two real-world scenarios, coming from energy distribution management and government publication services, which span numerous organizations and include long-tail aspects.

Figure 1-1 shows an interoperation hub being used to coordinate the activities of the six primary kinds of stakeholders and stakeholder organizations around hiring. Briefly, the kinds of organizations (and their associated services) that can participate in a collaboration supported by the hub are: Candidates, Human Resources (HR) Organization, Hiring Organizations, Evaluators, Travel Agency, and Reimbursement. There could be several hiring organizations (e.g., the Sales organization, the Logistics organization), each with its own workflows for managing the hiring process. We assume that the enterprise has a single HR organization responsible for hiring purposes. Participants and participating services can interact with the hub in several ways. For example, Hiring Organizations can post job descriptions, and candidates can apply for job openings and submit resumes and other materials. The hub does not prevent direct binary interactions; for example, a candidate may choose to interact directly with the designated Travel Agency, perhaps through their web-site, and create her itinerary. The hub can assist with such interactions, e.g., by recording the authorization for travel (perhaps from HR or perhaps from the Hiring Organization). The Travel Agency can access the hub for the travel authorization, and place a link to the itinerary. This enables the Hiring Organization and the Evaluators to access the itinerary when preparing for interviews. Finally, the Reimbursement organization can use the hub to access the airline and hotel invoices when processing the travel reimbursement request from the candidate. These interactions illustrate how an interoperation hub can facilitate information transfer between participant organizations, while giving them considerable autonomy and latitude with regards to how and when they provide the information or accomplish tasks.

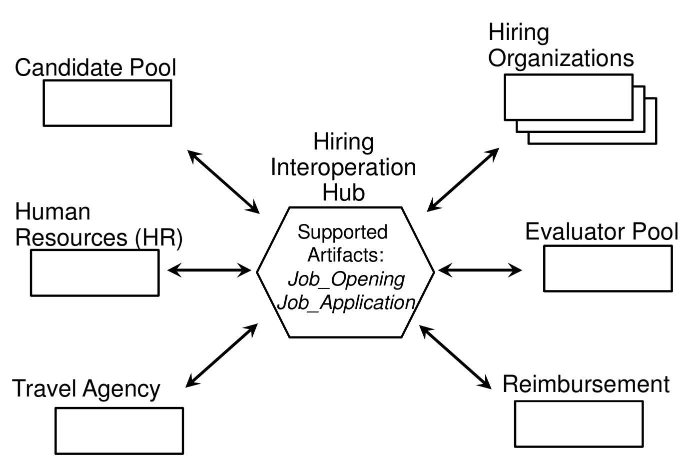

Figure 1-1: Example of an interoperation hub, which supports collaboration around hiring

As noted above, an interoperation hub differs from classical orchestrators in that it is largely *re-active*, and serves as a kind of *structured white board* to facilitate communication and collaboration between the different stakeholders. This is based on the premise that the participating services are self-motivated to contribute towards the common goals of the collaboration. Hubs can also be *pro-active*, in particular by notifying participating organizations of events or transitions that the organizations subscribe for.

Artifacts. The other central tenant of the ACSI research program is to build on the artifact-centric approach to modeling and implementing business operations and processes. This approach was originally introduced by IBM Research in 2003 [NiCa03], and has been developed since then at the methodology and systems levels by IBM [SNK+08]. Essential elements of the approach are also found in research works such as [RDtHI09, vdABEW01, BDW07, ABGM09], and have also been applied in various application domains, including government [P+04, PGGB09] and healthcare [BSB+08].

A *(dynamic) artifact* is a key conceptual entity that evolves as it moves through a business (or other) process. In the Hiring example, the interoperation hub maintains two artifact types, namely *Job Opening* and *Job Application*; Figure 1-2 illustrates the latter. As suggested there, an artifact type includes both a *data schema* and a *lifecycle schema*, which are tightly linked. As shown, the data schema has room for information about the candidate, the job openings he has applied to, the interviews, the trip information (if there are interview trips), the evaluations, and even the negotiations leading up to hiring. Importantly, the data schema

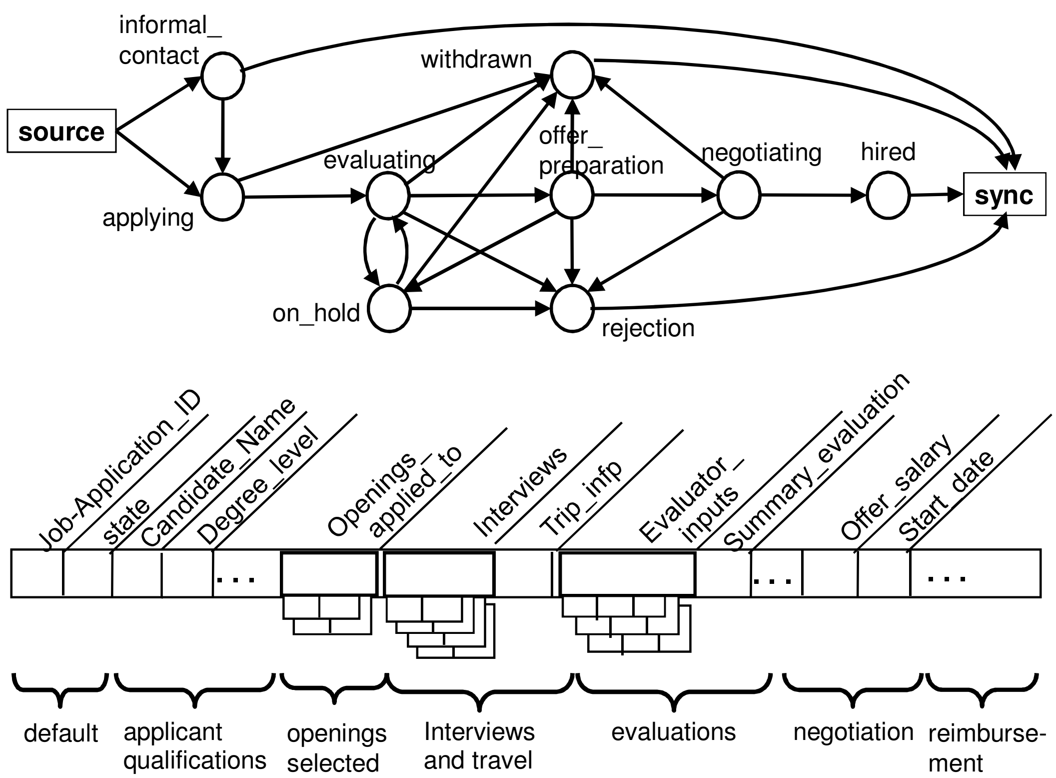

Figure 1-2: Job application artifact, showing both the data schema and the lifecycle schema

provides an end-to-end conceptual view of the key data for this artifact type; the data itself might be distributed in various ways between the participating organizations and the interoperation hub.

The lifecycle schema of an artifact type specifies the different ways that an artifact instance might evolve as it moves through the overall process. Much of the work on artifacts to date has used a form of finite state machines to specify these lifecycle schemas. (In ACSI, we will develop a more declarative approach to lifecycle specification.) As shown in the diagram, the typical lifecycle of a Job Application artifact instance is from application (perhaps with an informal contact first), followed by evaluation, and then possibly offer preparation, negotiation, and hiring. In the model used here, individual tasks (e.g., assign a ranking to a candidate, or perform an interview and record the results) occur within the states of the state machine. In the context of the interoperation hub, these tasks are typically performed by the various participating services (although the interoperation hub might perform some of these tasks itself). Also, it is typically the services that advance the lifecycle from one state to another when certain conditions, i.e., *guards*, on the artifact instance data are met.

Because they focus on the end-to-end lifecycle of key conceptual entities that pass through a business process, **artifact types can aid substantially in the management of "hand-offs" of data and processing between services and organizations**. For example, the Job Application artifact type makes it easy for stakeholders of the Travel Agency to understand how the data and processing they provide fits into the larger picture. Further, the artifact type provides a simple, intuitive structure for incorporating the data and processing performed by the Travel Agency into the larger context.

The artifact types used to specify the data and process management in an interoperation hub provide a **natural basis for specifying the access privileges of the different participating services**. For example, only HR, the Hiring Organization, and the candidate should be able to see an offered salary. Also, typically the candidate should not be aware of some of the states of the lifecycle schema, such as "on_hold", or "offer_preparation"; these should be collapsed into the state "evaluating" for the view presented to candidates. In [HNN09] initial ideas for systematic specification of such access controls are presented; these ideas will be developed further in the ACSI research.

Artifact-centric interoperation hubs and service collaborations in open service networks. The starting point of the ACSI research is to use artifacts to provide the underlying model of operations and processes in an interoperation hub, and thereby obtain the same benefits for the design, deployment, and maintenance of service collaborations that have been demonstrated by the use of artifacts in a single-enterprise setting. (Importantly, the services working with an interoperation hub are not required to be artifact-centric.)

Over the past 4 years, IBM Research has been developing methods and tools that use the artifact approach to design and deploy business operations and processes within a single enterprise, and has applied the technique in about two dozen engagements with IBM customers and internal IBM clients (see, e.g., [BCK+07, SNK+08, Ver09]). **The artifact approach has been shown to be effective in three profound ways:** (1) providing an end-to-end view of the operations that enables substantially improved communication between stakeholders, especially if they represent different "silo's" within the organization; (2) providing a unifying model of operations and processes, upon which process management, data management, linkage to legacy applications, and business rules can be layered; and (3) enabling a substantial speed-up over conventional techniques in the design and deployment of IT infrastructure to support business operations. **ACSI interoperation hubs, whose internal**

data and process management are based on artifacts, will provide the same advantages for service collaborations.

Another key advantage of the artifact-centric approach in the context of collaboration environments is that it **simplifies the modeling and management of so-called "points of variation"** [Ver09]. For example, in the Job Application artifact type of the Hiring example, a primary focus is on putting information into the data schema as the Job Application instance moves through its lifecycle. The lifecycle schema provides high-level guidelines about the sequencing for that information, but many details about how pieces of information are provided, and the particular sequence at which they arrive, are largely unspecified. The guards on state transitions can be customized to the different service collaborations (e.g., hiring for Indra vs. hiring for Collibra.) Furthermore, part of the ACSI research is to develop declarative approaches to specify lifecycle schemas that will enable the systematic modeling of richer forms of variation in interoperation hubs.

ACSI interoperation hubs will overcome many of the challenges current technologies face in the context of supporting open, long-tail service interoperation. Artifacts give a combined view of both the data and the processing that takes place in a service interoperation. While the participating services do not need to be artifact-centric themselves, the designers in charge of on-boarding and customization can take advantage of the global, end-to-end artifact-centric perspective, and their ability to support variations. Importantly, ACSI interoperation hubs will allow for incremental adoption: existing collaborations can migrate to full use of the hub in a series of steps, relying on the hub to gradually manage increasing amounts of their shared data and process interactions. In support of scalability, the hub concept spreads the work of adaptation to the participating services, and the ACSI research will develop auto-generation of adapters for some contexts. In support of long tails, the research will develop principled approaches to customization of the artifact types in a hub, and tools to enable the service owners to specify those customizations.

1.1.2 Scientific research stream

The driving goal of the ACSI research program is to build on the foundation of dynamic artifacts and interoperation hubs to develop scientific principles and techniques for a new, scalable approach to the effective and flexible design and management of collaboration environments in an open services network.

Towards this end, the scientific research stream consists of four tightly interrelated components (Figure 1-3). The first is the development of a formal, logic-based *ACSI Artifact Paradigm* that can support reasoning about, interacting with, and the evolution of artifact types, especially as they arise in interoperation hubs. The second is to develop the *ACSI Interoperation Hub Framework*, including the principles and techniques to support management of access privileges, a comprehensive design workbench, and an interoperation hub engine. (This work will be tightly inter-related with

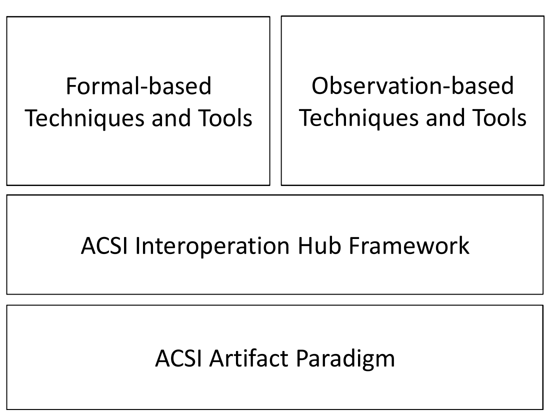

Figure 1-3: Components of scientific research stream

the technology research stream, which will prototype and integrate the techniques developed.) The other two components of the scientific research stream will provide additional key capabilities to make the artifact-centric interoperation hubs easy to design, participate in, and maintain. The *Formal-based Techniques and Tools* component will develop a scientific basis for managing views and evolution of artifact schemas, and develop verification algorithms and tools. The *Observation-based Techniques and Tools* component will adapt, generalize, and apply process mining and monitoring techniques to the context of artifact-centric interoperation hubs. The goals of the four components are now described. Further detail about the specific research techniques to be used, and how they will advance the current state of the art, is presented in Section 1.2 below.

The ACSI artifact paradigm. The research on the artifact paradigm is focused on achieving two of the seven targeted objectives of the ACSI research. While motivated by how the artifact approach can support interoperation hubs, these objectives will be achieved primarily through the development of the artifact paradigm itself.

Objective 2: Development and exploitation of declarative approaches for specifying the lifecycles of dynamic artifacts, to improve their applicability for use in interoperation hubs.

Objective 3: Development of a formal foundation for dynamic artifacts, including the development of a semantic layer, which will provide a principled basis for specifying and implementing views, transformation, and evolution of artifact schemas in the context of interoperation hubs.

To achieve Objective 2, the research will focus on forms of Event-Condition-Action (ECA) rules. This is because artifact instances, especially as they arise in interoperation hubs, are essentially reactive entities that take actions based on events generated by participating services or events in the external world. ECA rules can allow for lifecycles richer than finite state machines, for example, including parallelism and hierarchy, while still supporting a mathematically defined operational semantics. A declarative approach also holds the promise of enabling richer support for variations. The work of the PIs [BHS09, BGH+07, DHPV09, Hul08] will provide a starting point for this exploration. This work shall also draw inspiration from OMG's Semantics for Business Vocabulary and Rules (SBVR) standard [Obj08].

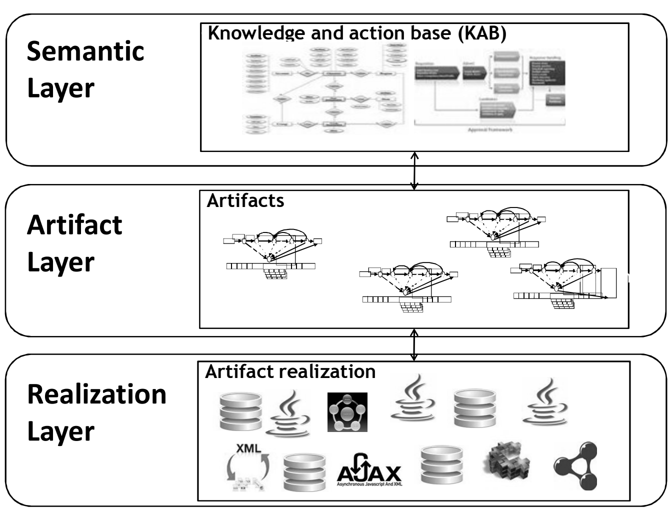

Figure 1-4: ACSI artifact paradigm

Achieving Objective 3 will require deeper scientific work, and will be achieved by developing a three-layered perspective on artifacts, as suggested in Figure 1-4.

The three layers of the ACSI artifact paradigm are as follows: The ***artifact layer***: in an instantiation of the paradigm, this consists of an artifact schema and a set of artifact instances that the system manages. The ***realization layer***: in an instantiation of the paradigm, this consists of the actual services for storing and managing data, executing and recording actions, and interfacing with the external world, all according to the artifact schema. The ***semantic layer***: in an instantiation of the paradigm, this consists of a conceptual specification of the static and dynamic aspects of the domain of interest. These are now described in more detail, along with key challenges that will be addressed in the research.

Artifact layer: This is the key layer of the ACSI artifact paradigm. It provides the actual artifact types that are available in the system, each, in turn, characterized by the manipulated data and the processes that manipulate them. Specifically, an *artifact type* in our approach is formed by the following elements: (i) an *artifact data schema*, which is a schema for holding the data relevant to the artifact type: (ii) a set of *artifact tasks* that can be invoked to modify the data in the artifact and to interact with the external world (I/O); and (iii) an *artifact lifecycle schema*, which specifies (either procedurally or declaratively) the permitted sequencings of artifact tasks. An *artifact schema* is a collection of artifact types.

The artifact data schema can be described as a data structure that contains all the information of interest for the end-to-end processing of the conceptual entities represented by the artifact type. The artifact tasks are the operations that are responsible of changing the data values in the artifact instances and of the interaction between the artifact instances with each other and the external world. The system may include a set of general artifact tasks, which can be used by several artifact types. The artifact lifecycle schema captures the process, expressed in terms of the artifact tasks and the artifact data schema that each artifact instance of this type is meant to carry out. Artifacts can communicate which each other through messages and shared tasks.

It is important to stress that **the languages used to specify artifact types are high-level but executable**. On the one hand, *artifact types are concrete enough* to make it possible to “compile” them into actual running programs, through suitable, model-driven techniques. On the other hand, *artifacts abstract from most implementation details*, providing a description of the activities of interest to the system at what we may call a “logical-level”, borrowing the terminology from relational databases. We will use this level of abstraction in most of the ACSI research, including work on verification and conformance testing.

Realization layer: This layer is based on an artifact engine for creating and processing artifact instances according to the guidelines of an artifact schema. As a starting point in the research we shall use the prototype Siena artifact system [CHP+08], developed by IBM Research, which implements a state-machine based form of artifacts. That system, like the other implementations of artifacts, assumes that all of the artifact instance data is managed by the engine, and essentially all of the task executions in the artifact types are performed by the engine. An important aspect of the research at the Realization Layer is to develop an approach that permits the artifact engine to use a hybrid approach, where services can choose whether to hold their own data (respectively, execute their own tasks) or let the hub hold the data (resp., execute some or all of the tasks).

Semantic layer: This fundamentally new component provides the conceptual representation of, and associated automated reasoning capabilities for, various aspects of interest to the system. This layer is centered around the new formal notion of **Knowledge and Action Base (KAB)**, that will hold information about both the static and dynamic aspects of artifact types and instances. In particular, a KAB will hold: (a) All the data of interest for the system; (b) Atomic actions, which correspond loosely to the tasks at the artifact layer; and (c) A family of dynamic constraints that cannot be violated when manipulating the data.

A KAB is the analog, for the context of transforming and integrating combinations of data and process, of the notion of knowledge base for the context of transforming and integrating across multiple, possibly diverse data sources. A KAB includes a knowledge base component, along with a tightly linked “action base” component. The knowledge base component consists of an ontology for data access and integration as proposed by the PIs in [CDGL+07b, CDGL+06, HMA+08, PLC+08, CDGL+09b, CDGL+09a], which is linked through suitable mappings to data sources, possibly maintained at participants’ sites [ZHK96, Len02, Hal03, CCDGL04, DGLLR07, Kol05, FKPT09]. Artifact data schemas can be seen as views over such a knowledge base component [Hal01, CDGLV07, CDGLR08]. The Action Base holds information about atomic actions, which include pre- and post-conditions concerning their impact, expressed in terms of the knowledge base; this in the spirit of AI research on reasoning about actions [Rei01, Sha97, San94]. The Action Base also holds information about temporal constraints (e.g., negotiations about hiring cannot begin until at least two interview evaluations have been obtained); these are in the spirit of situation calculus based high-level agent programming languages such as Golog, ConGolog and IndiGolog [LRL+97, DGLL00, SDGLL04], and especially in the spirit of temporal logics used for specification and verification of systems such as LTL, CTL, dynamic logics, mu-Calculus, etc. [EC80, Var96, Eme97, HKT00, CGP99, MP92, EF06].

The ACSI interoperation hub framework. This research component will develop the notion of interoperation hub along several dimensions, so that it can become a practical reality and reach its full potential as an enabler of open, scalable, long-tail service interoperation. This work will be essential for achieving the targeted Objective 1 mentioned in Subsection 1.1.1 above.

This work will be centered around creating a comprehensive *framework* for interoperation hubs and participating services, including (i) a capability for specifying and enforcing sophisticated access controls for different participating services, (ii) extensions of the interoperation hub notion of [HNN09] to support collaboration environments rather than single service collaborations; and (iii) the development of techniques so that the Hub Engine can support data and tasks in both virtual and materialized ways.

The ACSI interoperation hub framework will be designed to take full advantage of the richness of the ACSI artifact paradigm, and participants can select increasingly sophisticated interoperation modalities. The goal is to **simplify the job, of both service collaboration designers and the users of participating services, to participate in, customize, or substantially re-design service collaborations**. We foresee essentially three forms of interoperation, each with increasing flexibility: (a) ***Use-only interoperation***: Making use of one or more artifact types provided by the interoperation hub; (b) ***Customize-only interoperation***: Customizing one or more artifact types in the interoperation hub, thus introducing new, but derived, artifact types in the hub. The customization essentially amounts to using a subset of the full capabilities of the artifact type. (c) ***Full interoperation***: Creating new artifact types in the interoperation hub.

To the increasing level of flexibility obviously corresponds an increasing level of reasoning support required by the interoperation hub in setting up the new artifacts and to manage them at runtime.

Formal-based techniques and tools. This research component will include primarily two bodies of work. The first is the development of techniques for reasoning about and using the Knowledge and Action Base (KAB) that form the semantic layer of the ACSI artifact paradigm (part of Objective 3 above). The second is the development of techniques for *verification* and *synthesis* of artifact type designs (Objective 4 below).

In connection with KABs, the research work in the ACSI artifact paradigm will establish the basic framework, and the research in the Formal-based Techniques and Tools will focus on reasoning about KABs, mappings between a KAB and artifact schemas, and a theory that enables the comparison and mapping between artifact schemas by referring to a common KAB. (This latter application of the KAB is analogous to the use of ontologies to serve as a reference model against which two or more database schemas can be compared. It is also analogous to the main product suite of partner Collibra. In particular, Business Semantics Glossary supports the specification of a semantic data layer; using SBVR or e.g., in the NIAM data model. The Business Semantics Studio enables the specification of mappings from one database schema through the semantic data layer to a second database schema. Finally, Business Semantics Enabler automatically generates data services based on the mappings that transform data from one schema to another.) Most of the work on the KABs will fall into one of three categories. The first is extensions of the theory of ontologies for data access and integration to support update of the extensional data, at least in the context of the kinds of update that arise in artifact-centric interoperation hubs [LLMW06, DGLPR07, DGLPR09]. The second is to decide on the basic model for the Action Base and develop reasoning capabilities for verifying dynamic properties of how the information in the knowledge base component changes (e.g., reachability of one state from another via an unbounded sequence of atomic actions) [CDGLR07, BBL09, HBHP09, CL08, SDG09]. Basic mechanisms for querying KABs form a key component of the first and second categories [CDGL98, CDGL08, CDGL+07a, CDGL+07b, GHLS07, OCE08, LTW09]. The third category is to study the possible relationships between KABs and artifact schemas, including support for conformance testing and evolution. This work will include the prototype implementation of targeted tools that embody reasoning and artifact schema evolution capabilities based on KABs.

The second portion of the Formal-based research will address the fourth targeted research objective of ACSI.

Objective 4: Development of techniques and tools to support automatic verification of interoperation hub properties at both design- and run- time, and partial synthesis of interoperation hub components.

This topic is deeply challenging, because artifacts involve both data and process, which means that the underlying state space is essentially infinite. This forces us to develop new extensions for traditional model-checking based approaches to verification. Preliminary work by the PIs [BGH+07, DHPV09], that provides some first verification results for artifacts, shows that quite restrictive abstractions are needed to obtain contexts in which useful properties are decidable. Thus, a key element of the work will be to identify abstract settings that (a) enable specification of properties that are of interest in the context of interoperation hubs, yet (b) permit effective verification in practice. This work will focus on extensions to known model-checking techniques [CGP99], and using logics other than LTL (e.g., epistemic logics [FHMV95]), deontic logic and ATL [AHK02]), to specify such abstractions. The goal is to enable design-time verification of key properties about interoperation hubs and services, such as compliance of a service to a hub, whether a hub will enforce binary *commitments* [DMCS05] or contracts between services, and whether a service collaboration will achieve specified objectives. The research shall implement a verification toolkit based on some of the techniques developed, for verification in the artifact and semantic layers. To maximize reuse, the toolkit may build upon existing model checkers such as NuSMV [C+02] or MCMAS

[LQR09] by adapting and extending the verification algorithms to the case of artifact centric systems.

While the design-time verification toolkit produced in the project will enable us to predict the behaviour of artifact-based environments we also wish to be able to monitor at run-time particular executions of artifact based environments. Specifically, we would like to be able to detect whether particular properties of interest, e.g., violations of particular conditions happen at run-time, and, if so, whether any recovery mechanism implemented is capable of restoring the system to correct functioning behaviour following these events. While run-time monitoring for systems and for web services is obviously a very active area of research [RV], the area has so far focussed on processes with finite states. In the project we wish to monitor artifact-based environments, so we will need to be able to reason about processes with infinite data. This leads to consider as a starting point approaches where a "symbolic" treatment of the state-space is employed as in [LPSS09,RSE08]. In the project we aim to identify a suitable symbolic technique and extend it by using automatic abstraction techniques to enable the verification of both data and processes in the artifact model.

A final portion of this research is to develop algorithms (along with prototype implementation) for the automatic synthesis and composition of artifacts. As with analysis, this will be an advance over the state of the art because of the prominent role that data has in artifacts. The goal of this portion of the work is to enable high-level specification of the goals that an interoperation hub should satisfy, and then automatically create the artifact schema for the hub, along with appropriate access controls for the different participating services. This research will build on the PIs' work [FHS09], which gives a synthesis algorithm for artifact types in a restricted setting, and also the PIs' work on service and agent composition [CDGL+08b, BCDG+05]. The research will also generalize techniques of doing synthesis by starting with a template and then customizing it to satisfy desired goals, in the spirit of [SPM09].

Observation-based techniques and tools. This stream will develop fundamentally new approaches to process mining. The overall objective here is:

Objective 5: Development of techniques and tools for process mining of services and artifact-centric interoperation hubs, to support on-line compliance checking, auto-generation of adapters, and the discovery of artifact-centric models that underlie a service's behaviour.

At its core, this part of the ACSI research will exploit a striking relationship between process mining and the artifact approach. The starting premise in most process mining work is to examine logs of the execution of (large-scale) business process engines, correlate the events from those logs in useful ways, and from there infer the behavior patterns that are being followed, typically as activity flows. In existing process mining techniques (e.g. those implemented in the ProM framework [vdV+99]) the log events are correlated based on a key (also known as case id), to create what are called *process instance traces*. A process instance trace groups the sequence of events that occurred to a single conceptual entity and the data associated with these events. For example, in the context of a hiring process, one might build process instance traces based on Candidates, on Reviewers, on Job Applications, or on Job Openings, etc. In many cases, based on the type of conceptual entity chosen to correlate log events, the resulting process instance trace will correspond to an artifact instance, which may be either implicit or explicit in the design of the business process being analyzed.

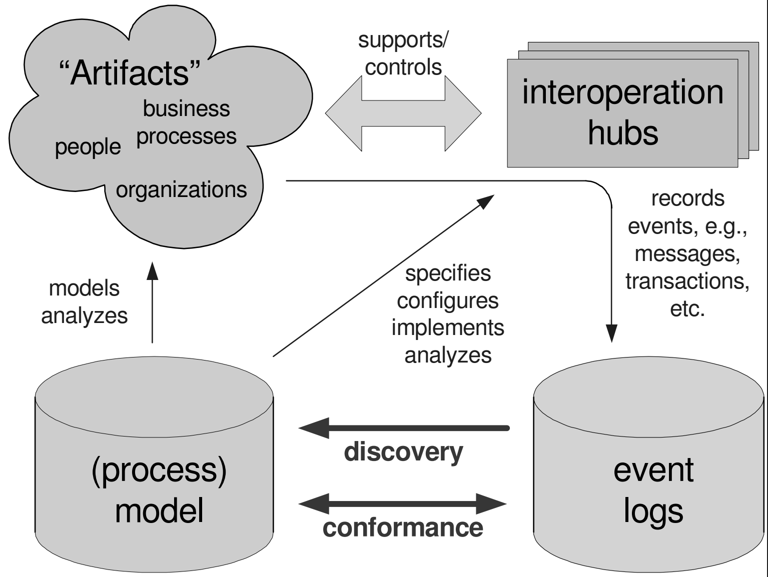

Figure 1-5: The relation between process mining, artifacts, and interoperation hubs

Figure 1-5 shows the relationship between process mining, artifacts, and information hubs. First of all, the artifact paradigm much better reflects real-life processes than the traditional monolithic workflow paradigm. Various case studies have demonstrated this, i.e., rather than finding large process instances, smaller interrelated artifacts are discovered. Secondly, the interoperation hub will generate the ability to collect high-quality data. This can be used for conformance and performance analysis. Moreover, when multiple organizations use the hub, cross organizational mining can be performed.

A challenge in ACSI is to automatically discover artifact types, their life-cycles, and the relations them. This is far from trivial as existing process mining techniques consider only one artifact type in isolation. At the same time it is very important, as real life processes emerge from the interaction between various articfact (types). This connection between process mining and artifacts will serve as a stepping stone in the research to be performed to achieve this objective.

Process mining research will be extended in ACSI in two substantial ways: a shift to **on-line analysis**, and a shift to **incoporate data more explicitly** as it arises in the artifact approach. The area of process mining has focused to date primarily on the *post mortem* analysis. For the context of large-scale collaboration environments, however, on-line analogs are required. Turning to data, existing process mining techniques are focused on identifying control-flow dependencies [vdAvDH+03, vdAWM04]. No comprehensive technique exists for mining data dependencies and the discovery of artifacts. Current techniques cannot discover which services read and write which attributes of the artifacts and at which point in time.

The extended process mining techniques will be applied in primarily three ways during the ACSI research. The first is **conformance checking, both of an interoperation hub's actual behavior against its specified behavior, and of a service with the behavior required by an interoperation hub.** The purpose of conformance checking [vDADO+08] is to ascertain the conformance of the execution log of an existing system with respect to a given artifact-centric model and to pinpoint deviations between the logs and the model (a.k.a. gap analysis). Conformance checking is a central monitoring activity in service interoperation hubs, since deviations between the actual hub's behavior and the model are indicators of interoperability problems. Also, the conformance of services against a hub can help to determine whether a service can participate in a service collaboration. Moreover, mismatches identified by the conformance checker will facilitate **on-the-fly synthesis of adapters** for resolving those mismatches. This will be a key technology for simplifying the on-boarding of services into service collaborations.

Finally, the ACSI research will develop approaches for **artifact discovery;** that is, constructing an artifact schema from a set of execution logs. In this project, we position artifact discovery as a central maintenance activity in interoperation hubs. By constructing operational models of how a hub actually behaves from service interaction logs, we will be able to reason about the evolution of the hub and the changes that occur in the hub over time. In particular, we will be able to pinpoint how the processes supported by a hub are affected by the arrival or departure of services, and we will be able to monitor data quality.

1.1.3 Technology research stream

This stream of the research is focused on the following objective.

Objective 6: Design and development of the ACSI Hub System, a substantial prototype realization of an integrated ACSI interoperation hub framework, including a Design Workbench and Hub Engine, that incorporates the targeted scientific advances achieved in the ACSI research in support of access controls, process mining, auto-generation, verification, synthesis, and evolution.

Figure 1-6 shows the conceptual architecture of the integrated ACSI Hub System that will be created over the three years of the project. This includes prototype capabilities from all four of the components of the Scientific Research Stream, and also new code for the integration. As a starting point, this stream will use the Siena artifact system [CHP+08], developed by IBM Research and now also being used for research at partner University of Rome La Sapienza. The Technology Stream will begin in the first year, and will be highly useful in helping to ensure throughout the life of the ACSI project that the Scientific Stream addresses the critical problems, and does not leave substantial gaps in terms of framework or tooling. The architecture is centered around two primary components, the *Design Workbench* and the *(Interoperation) Hub Engine*, shown here as top-level solid-line rectangles on the left. The Hub Engine interacts with the Participating Services, shown in the lower right.

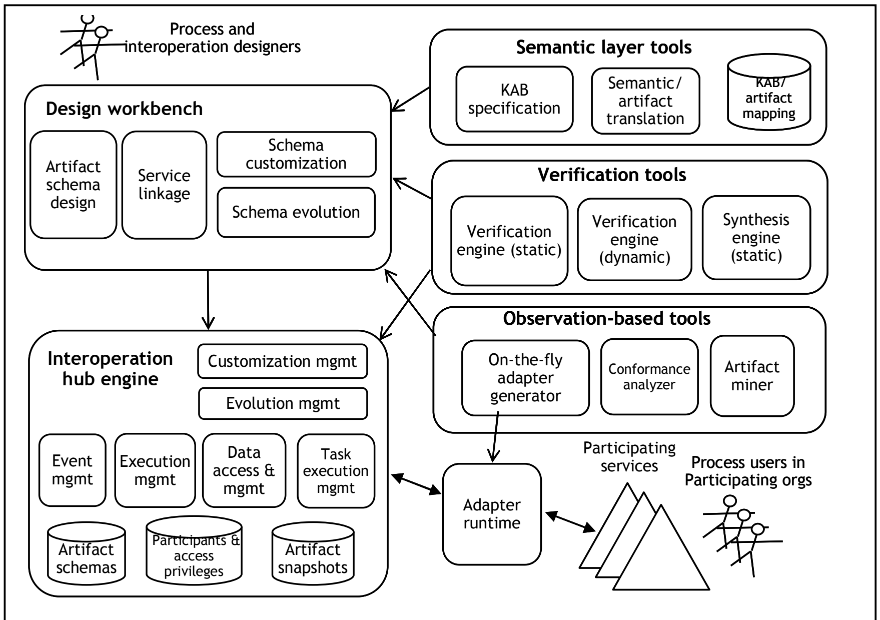

Figure 1-6: Conceptual architecture ACSI Hub System

The **Hub Engine** provides the run-time support for an interoperation hub. When operating, it maintains three kinds of information (shown using "disk icons") on: (a) the relevant artifact schemas; (b) the participants and their access rights; and also (c) the artifact snapshots, that is, for each active artifact instance, the Hub Engine will maintain a record of its current state, including the values held by all of the attributes in the artifact data schema. Normal operation is achieved by the four components shown in the middle layer of the Hub Engine. The *Event Management* component maintains a queue for incoming events, and associates them with the appropriate artifact instances. It also handles sending messages going to the participating services. The *Execution Management* component looks at each incoming event, refers to the relevant artifact instance and artifact type, and then determines what should be done. Typically this means that a task will be invoked. This component does not perform "actual work", rather, it dictates what work should be done, and delegates it appropriately. The *Data Access/Mgmt* component will make data associated to a given artifact instance available, even though the data might be stored within the Hub Engine or in a participating service. Similarly, the *Task Execution* component can manage getting a task accomplished either locally or by a participating service. Finally the *Customization Management* and *Evolution Management* are self-explanatory; these will also manage "in-flight" artifact instances during evolution. The Hub Engine will be integrated with the dynamic verification engine.

The **Design workbench** has four main components, which support designing schemas, linking services to deployed schemas, customizing schemas, and evolving schemas. Each of these will take advantage of auxiliary capabilities shown on the right side of the diagram as they become available in the ACSI research.

The **Semantic layer tools** component enables setting up a KAB, and then using it in connection with one or more artifact types, in particular to support artifact schema evolution. The **Verification Tools** component holds the static and dynamic verification engines, along with the synthesis engine. The **Observation-based Tools** component includes the *Conformance Analyzer*, the *On-the-fly Adapter Generator*, that is used to generate run-time *Adapters*; and the *Artifact Miner*.

1.1.4 Validation stream

The objective of this stream is:

Objective 7: Demonstrate and measure the effectiveness of the ACSI research results and ACSI Hub System on two use cases that involve service collaborations in an open service network, and exhibit long-tail characteristics.

We first recall the four metrics against which the ACSI technology will be evaluated and then outline the two use cases. The methodology to be used is described in WP5. We plan to incorporate the envisioned set of semantic patterns into the SEMIC.EU platform, so that they can be adopted by third parties. Also we will publish case studies on the e-Practice platform. Both platforms are owned by the former IDABC (soon ISA) programme of DIGIT.

The high potential impact of the ACSI System Hub can best be demonstrated through validation of various use cases. Therefore, we will strive to establish at a very early stage in the project, a Special Advisory Board that will, among other things, provide additional use cases for validation (refer to section 3.2.2).

Recall from the beginning of Subsection 1.1 that the four objective measures of the success of the ACSI technology are, briefly: (1) at least 40% reduction in design and deployment costs for service collaborations; (2) at least 20% reduction in costs of on-boarding and maintenance; (3) at least 30% reduction in on-going manual activity to support service collaborations, and (4) at least 90% of the data transformation will be automated. In metric (1), the percentage reduction is based on an experience of using the artifact-centric approach in a single-enterprise setting on a large-scale business transformation project within IBM. Additional savings may accrue for cases where a single ACSI interoperation hub is supporting multiple service collaborations. Metric (2) is rather conservative, because we do not have any direct experience with something analogous to maintenance of interoperation hubs. Our expectation is that the on-boarding costs will be reduced because of the intuitive nature of artifact schemas and the on-the-fly adapter generation, and that maintenance will be simpler because of the use of artifacts. Metric (4) is based on preliminary investigations of the two use cases, and also several other examples from local and regional government and also from upstream supply chains. Metric (3) relies in part on Metric (4), but is otherwise conservative, again because of a lack of direct experience with this characteristic

The research approach in this project is a combination of theoretical work, computational implementation and testing, and field validation. In parallel to the theoretical work two use cases will be defined and validated in the fields of Energy distribution and European Journal. For each of the case studies we will identify scenarios in which the project results can be assessed according to the predefined metrics. At the end, the accumulation of the validation efforts should allow triangulation across the scenarios tested.

In order to evaluate the effectiveness of applying the ACSI Hub System, the four defined metrics will be assessed and evaluated before and after in both scenarios.

Energy distribution management case study: This use case is focused on the data gathering and billing management system within the whole process of electrical energy production and distribution. Most big energy suppliers in Spain depend on large numbers of small providers (usually SMEs) that use different measurement devices, different measurement timings, different databases, etc. Nevertheless, the energy supplier must maintain an integrated view of this data, and use it for several mission critical processes, that combine and analyze measurement data from various sources, and distribute the information for publication and/or further processing by other components, such as the billing system. Electricity producing enterprises make sales offers of a certain amount of electricity at a certain price for a time of the day. Continuing measurement is needed to compare energy consumed and energy billed, both to support auctions and to check for inconsistencies. The typical stakeholders are System Operator, in charge of supply continuity and security; The Electric System Commission, watching for the interests of the consumers; European Union, which establishes the general frame of the electric system for all the countries of the EU; and suppliers, which nowadays can be SMEs, particulars and energy producer companies. In most of the operations, several systems take part, interchanging information in many different ways using different technologies and techniques. The pilot will address a full cycle of beginning to end processing using a limited set of data, but heterogeneous enough to be close to real life, to test both the technology and methodology developed in ACSI.

European journal case study: The Publications Office of the European Union (OPOCE) is an inter-institutional office whose task is to publish the publications of the institutions of the European Communities and the European Union. OPOCE publishes daily the **Official Journal of the European Union** in 22 or 23 languages (in the Irish language when the publication is necessary), a unique phenomenon in the publishing world. It publishes or co-publishes publications entrusted to the Office in the context of the institutions' communication activities. Moreover, OPOCE offers a number of *online services* giving free access to information on EU law **(EUR-Lex)**, EU publications **(EU Bookshop)**, public procurement **(TED)**, and EU research and development **(CORDIS)**. Their figures indicate the importance of this office producing more than one million pages every year with about 600 staff members [Eur].

In order to create a new online service – which corresponds to creating a new service collaboration to be supported by an interoperation hub -- the publication office collects data from different types of stakeholders, and compiles this data into a uniform and presentable format. Currently most of the activities are completely manual, in part because of the variability between the different online services. There is a lack of semantic interoperability (i.e., the ability to use differing data which is structured in different ways) among the stakeholders of the business process around online service delivery, which contributes heavily to the amount of manual effort. The variability is exacerbated by the many cultural and regional differences between stakeholders.

1.1.5 Compliance to objective ICT-2009.1.2 (b)

The ACSI project addresses strategic Objective ICT-2009.1.2: Internet of Services, Software and Virtualisation. b) Highly Innovative Service / Software Engineering. In the table below we list the specific target outcomes as they appear in the call, and describe how ACSI complies with these outcomes.

Specific EC Objectives	ACSI Compliance
Service / Software engineering methods and tools *covering automatic support at run-time for decisions and changes that are currently adopted at design time.*	The research will support this objective in two ways: First, by leveraging extensions of verification and process mining, techniques and tools will be developed to support some automation within the three year research activity. Second, the research will establish a new, formal foundation for artifacts and interoperation hubs that can provide the basis for automation of tasks such as hub evolution in the future.
Focus is on innovative approaches to very large, dynamic open service network,	Orchestration is founded on the premise of knowing in advance what services will participate in the interoperation environment. In contrast, an interoperation hub essentially establishes a *de facto* standard that many separate service collaborations can take advantage of. This provides a level of scalability not achieved by orchestration or choreography. Also, the artifact-centric approach, especially when it is extended to have declarative lifecycles, will provide rich support for variations. This is a second feature of ACSI interoperation hubs that enables scalability.
user development of services/software, systems evolvability and acquisition,	This is addressed at a fundamental level, by the use of dynamic artifacts, which provide an holistic and intuitive marriage of data and process. Verification tools will also help users as they design hubs and on-board services. Process mining techniques will be used to simplify discovery and conformance testing, and to auto-generate adapters. Finally, the semantic layer of the artifact paradigm provides a principled basis for specifying hub evolution. The Design Workbench will tie these capabilities together.
reasoning and incorporation of domain knowledge in all phases of the service/software lifecycle	This is supported in two complimentary ways. First, the verification and process mining work both enable reasoning about (future or past) enactments of the business processes supported by an interoperation hub. They will both address the fact that artifacts hold data; this typically includes a fair amount of relevant information from the surrounding aspects. Second, the semantic layer in the Artifact Paradigm is based on the new notion of Knowledge and Action Base (KAB), and will support reasoning.

Specific EC Objectives	ACSI Compliance
High-level *description and executable* ***languages*** *for services/software*	The artifacts themselves form the basis for a high-level, highly intuitive language for specifying the data and processing aspects of the interoperating services. The artifact specifications can be implemented using the model-driven style. Artifacts have been shown at a pragmatic level to provide a substantial advantage over current approaches for specifying services and business processes. By investigating the foundations of the artifact paradigm and developing the semantic layer, the research will improve on and extend the reach of current artifact specification languages. The research will also develop new declarative approaches for specifying artifacts.
Verification and validation methods*, tools and techniques assuring the quality of open, large-scale, dynamic service systems without fixed system boundaries…*	The high-level nature of artifacts allows in principle for automated verification. The research will address the presence of data as a core component in artifacts, leading to special kinds of infinite state systems. The verification questions investigated will be motivated from the use of artifacts in interoperation hubs, and will thus address the issues of open, large-scale, dynamic, and lack of fixed boundaries. The research will develop verification techniques for design-time, deploy-time, and run-time.
Methods, tools and approaches specifically supporting the development, deployment and evolution of ***open source software****…*	The research makes indirect contributions towards this goal along two dimensions. We intend that much of the software developed in the research will itself be open-source. More generally, the interoperation hub framework provides a natural context within which open source developers can create highly visible, widely used software that can have substantial impact.

B1.2 Progress beyond the state of the art

We describe now the current state-of-the-art in those areas that will most be impacted by ACSI, and the concrete advancements that ACSI will bring to each of them.

1.2.1 Business process management and artifact paradigm

Current approaches to BPM, while primarily centered around activity flows and supporting design by process decomposition, often do provide streams of activity around information modeling, integration of legacy applications, and business rules and requirements. These are all helpful in providing an intuitively accessible, business-level view, at varying levels of abstraction, into different dimensions of an underlying family of operations. However, the underlying conceptual models for these different elements are disparate: each is based on a different core focus, each has different styles of problem decomposition. As a result, the integration of these diverse streams in the IT-level implementation of a BPM design is often ad hoc, which adds substantial complexity when trying to evolve the implementation to meet new business goals, and when trying to interoperate between silos or enterprises whose underlying BPM designs were created separately. Further, there is no broadly accepted and principled approach for describing how the streams of a BPM design work together, and more to the point, there is no such approach for describing how a family of BPM-based services can collaborate together.

As described in Section 1.1, the dynamic artifact approach [NC03, SNK+08] helps to overcome the disparity of conceptual models used in BPM, by providing an holistic model for specifying the end-to-end lifecycle, and associated data, of key conceptual entities that move through the business operations. The work on artifacts has focused primarily on methodology and systems issues, and the single-enterprise context.

At a very high level, dynamic artifacts have goals similar to both business objects (e.g., [EEO+98]) and agents as applied to BPM (e.g., [WJK00]). In particular, all three approaches focus on a handful of key constructs that can be used to specify portions of business operations, and that can interact in order to specify a large or full set of operations of an enterprise. But there are fundamental differences at two levels, one concerning the underlying philosophic motivations, and the other concerning the specific constructs used.

At the philosophic level, business objects are intended to be re-usable across several or many enterprises, and focus on relatively small-grained collections of functionality. Also, as suggested in [EEO+98], it is common to consider three categories of business objects, for entities, for processes, and for events. In contrast, dynamic artifact types are intended to provide an end-to-end view of the life of a key conceptual entity in an enterprise. BPM agents, at the philosophic level, are generally focused on clusters of activities that are performed by the sub-organizations of an enterprise. For example, in the agent-based design presented in [WJK00], each agent corresponds closely to the activities of a division of the company (e.g., network design, legal, customer vetting). This contrasts starkly with the dynamic artifact approach, which would typically focus on an artifact that will clearly cut across multiple sub-organizations and silos of the enterprise.

At the level of constructs used, both business objects and BPM agents find their roots in object-oriented programming. Both focus on encapsulating data and processing, and focus on exposing a family of services. Interaction is specified in terms of message-passing protocols. Business objects typically do not include specification of the overall lifecycle of their instances. In contrast, dynamic artifacts are based fundamentally on the marriage of an information schema and a lifecycle schema, both considered as first-class citizens having equal prominence. While there may be access controls over how, when, and to whom the data from the information model is revealed, that data is directly visible and accessible to the various agents (human or automated) involved with the artifact type. Together, the direct access to the data managed by the artifact type and the structuring of processing according to the lifecycle schema enable a perspective that is more declarative than afforded by either business objects or BPM agents. Another main differentiator between artifacts and agents is that agents are by definition proactive, whereas artifacts are passive entities. This apparently simple difference has profound implications in the underlying architectures.

In terms of modelling and execution, although several approaches to modelling and executing business processes based on (business) objects have been developed, (e.g. OCoN [WWG01], BML [WPD+03]) and Agent-Object-Relationship Modeling Language (AORML) [Wag03], these paradigms are focused on modelling and are not backed by an execution environment as envisaged in the ACSI hub platform. They also do not tie up with the problem of enabling interoperability between services in heterogeneous environments.

<u>Progress in ACSI beyond the state of the art</u>: The ACSI research will extend on the dynamic artifact approach in three key ways: (a) development of declarative approaches [MPvdA+10, PvdA06, vdAAtH+09, vdAP06] to specify artifact lifecycles, (b) development of a formal foundation for artifacts including a semantic layer to enable a principled approach to artifact schema evolution, and (c) the generalization of artifacts to support service collaborations involving multiple enterprises (discussed in Subsection 1.2.2.). The research towards (a) will build on the PIs' research on using ECA-style languages to study formal aspects of artifacts [BHS09, BGH+07, DHPV09, Hul08], but here with a focus on developing a pragmatic rather than theoretical language. Goals for the language include support for parallelism, hierarchy, and more explicit ways of supporting variation. As described in some detail in Subsection 1.1.2, the research for (b) will develop the ACSI artifact paradigm, with three layers: semantic, artifact, and realization. (See Subsections 1.2.3 and 1.2.4)

1.2.2 Service interoperation

There are today two dominant approaches to supporting service interoperation, namely, orchestration and choreography [ACKM04, Su08]. Both are based on the Service Oriented Architecture (SOA), which gives prominence to message-based interfaces to services. As a result, both tend to obscure the data of common interest to multiple services in a service collaboration. Also, neither approach is equipped, at a fundamental level, to handle variation across service collaborations in a principled manner. With orchestration, the design of the orchestrator is typically labor-intensive, and so the paradigm does not scale well in supporting many service collaborations. With choreography [BF08, KL08], the premise is to establish a comprehensive family of constraints on how a cluster of services should interact to achieve common goals. This, too, does not scale well to support minor variations between how different service collaborations will achieve their goals.

<u>Progress in ACSI beyond the state of the art</u>: As discussed in more depth in Subsection 1.1.1, the ACSI research will substantially develop the notion of artifact-centric interoperation hub, which was recently introduced by the PIs in a preliminary work [HNN09]. The interoperation hub notion borrows from both orchestration and choreography, but provides a new point in the spectrum of interoperation support possibilities. By using dynamic artifacts to specify the internal logic of ACSI interoperation hubs, designers and participating services can take advantage of an end-to-end view of service collaborations, and understand both the key data and the lifecycle of the entities that pass through it. The ACSI research will substantially fill out the notion of artifact-centric interoperation hub, including the development of a framework for specifying and enforcing access controls, the generalization of interoperation hubs to support environments with numerous service collaborations, the design and prototyping of a Design Workbench, and the prototyping of a Hub Engine that can operate as an SaaS. The Design Workbench will include access to the outputs of the Scientific Stream of the ACSI research. The Hub Engine will use the Siena artifact system [CHP+08] as the starting point, but will make substantial extensions, including adaptation to permit data to be stored, and tasks performed, either locally in the engine or remotely in the participating services.

1.2.3 Knowledge and Action Base – Focus on data and knowledge

The knowledge base portion of the KBA will be based on ontology-languages for data access and integration. *Ontology-based data access* is the problem of accessing concrete data through the conceptual level offered by an ontology. *Ontology-based data integration* is a form of "virtual" data integration in which the global view over all data of interest is an ontology with full querying support, the data remain at the sources which are heterogeneous and independent, and suitable mapping mechanisms are used to specify how to extract the data from the sources to answer queries over the ontology. The starting point of the work in ACSI will be the work on specific description logics for representing ontologies for which query answering has essentially the same cost as in relational databases [CDGL+05, CDGL+07]. These logics are ideal for ontology-based data integration, since they allow for reformulating queries taking into account both the ontology and the mappings to the data sources [PLC+08, CDGL+09]. Also, we will leverage powerful query languages studied in the context of ontologies, which go from conjunctive queries to first-order logic epistemic queries [CDGL+07a], and we will take into consideration the recent literature on ontology module extraction [CGHKS07, KPS+09]. So-called "write-also" data integration, where the objective is to perform updates on data sources through the ontology providing the global view, is a largely unexplored problem that must be addressed in the ACSI context. We will use as a starting point on the one hand traditional work on view-update in databases [AHV95] and recent literature on inverse mappings in data exchange [Fag07, FKPT08], and on the other hand the initial research on instance-level updates in ontologies [DGLPR06, LLMW06, DGLPR07, DGLPR09].

Progress in ACSI beyond the state of the art: We will make use of ontology-languages optimized for ontology-based data access to devise *powerful mechanisms for defining views over ontologies* that enable to circumscribe in a suitable way the information accessed, as required to capture artifact data schemas. In doing this, we will rely on the existing powerful query languages for ontologies and on the available module extraction techniques. However, we will extend and adapt current techniques to take into account that the view mechanism needed to capture artifacts is targeted towards individual instances, in contrast to typical view mechanisms in databases that rely on queries collecting entire sets of instances. The *theory of views over ontologies* for data access to be developed in ACSI will provide a strong formal basis for understanding rigorously relationships among the data components of different artifacts. The research will combine this with *ontology-based data integration* over views, since much of the data in the interoperation hub will in fact be virtual, i.e., held by the participating services. Artifacts will update their data component during their lifecycle. From the semantic-layer point of view these updates will be data (or instance-level) updates over an ontology. Note that by combining ontology update with the use of ontologies for data integration, we open up for *write-also data integration*, a very prominent issue that will be studied in ACSI. Tools implementing selected techniques for reasoning over views and view management, and for ontology-based read-only and write-also data integration in the context of the ACSI artifact paradigm will be realized.

1.2.4 Knowledge and Action Base – Focus on actions

For representing the Action Base portion of KABs, good starting points are action theories developed in Reasoning about Actions in AI [Rei01, Sha97, San94], combined with high-level agent programs such as ConGolog [LRL+97, DGLL00, SDGLL04], FOL-variants of temporal logics used in verification such as LTL, CTL, dynamic logics, or Mu-Calculus [EC80, Var96, Eme97, HKT00, CGP99, MP92, EF06]. We stress that for the Action Base component, we will take an approach that is radically different form the recent work on ontologies for processes such as OWL-S <http://www.w3.org/Submission/OWL-S/> and WSMO <http://www.wsmo.org/>, since our approach is strongly tailored towards representing and reasoning the dynamics of processes, exactly as in the Verification literature, and not on their static description as in the Ontology literature. Another approach that has been advocated is to use dynamic rules (e.g., (E)CA rules or SBVR rules) [dB09], possibly integrated with some temporal logic formalism. In any case, we are confronted with the necessity of modeling how the state of the ontology and the associated data evolve during the execution of a process. Unfortunately this problem is notoriously difficult: even for very simple combinations of the static and dynamic aspects: inference over the temporal evolution of an ontology is undecidable [WZ99]. Recently there has been some progress in studying how to restrict on the one hand the expressive power of both the temporal formalism and the ontology language, and on the other hand the interaction between the static and dynamic aspects, to achieve decidability [AT08, AKRZ09]. But the limitations in expressive power are rather significant, and decidability turns out to lack robustness.

Progress in ACSI beyond the state of the art: Marrying the description of the static aspects of the domain of interest (the data) with the description of its dynamic aspects (the processes) is the key element for actually realizing the activity-base at the semantic layer. However, considering the essentially negative results on decidability obtained so far, in ACSI we propose to tackle this problem following a different approach. We will consider the ontology in the Action Base as a knowledge base under Levesque's functional view [Lev84], i.e., a sort of rich data type constituted by the information in the ontology with the associated data, and two basic kinds of operations: ASK (based on query answering) and TELL (based on instance-level update). This view greatly simplifies the problem of combining static aspects with dynamic ones: enabling to base the first ones on logical implication and the second ones on model checking. We stress that such a model checking could involve *infinite states* in the case of unbounded processes, so some forms of *abstraction* will be required. The project will pursue this challenge through different research directions: *(i)* exploit promising work based on the *functional views of ontologies* developed within the Description Logics community [CDGLR07, BGL08, BBL09, HBHP09]; *(ii)* exploit very recent work from reasoning about actions in AI, which looks into *verification of (infinite state)* ConGolog *programs* based on Situation Calculus [DGLL00], through a notion of "characteristic graph" [CL08, CL09, SDG09]; *(iii)* exploit the *theory of conjunctive queries* in relational databases, which is the basis of most recent results on data exchange [FKP05, FKPT09] and notions like "weak acyclicity" that guarantees that unbounded application of tuple-generating dependencies (that could be the basis for action specification) will converge to finite resulting states; *(iv)* exploit the work on *abstraction on data* developed in databases and recently at the base of very fruitful work on mixing processes and data in various contexts [BCDG+05, DSV07, Via09, DHPV09, CDGHS09]. A key driver for this research will be to enable useful support for the evolution of artifact schemas in the context of service interoperation. A KAB reasoning toolkit will be developed based on the framework and algorithms obtained. It is expected that the toolkit will be released as open source for maximum impact and potential reuse and extension by the community.

1.2.5 Verification in the artifact paradigm

The agent paradigm [Woo09] is central to the formalization and implementation of autonomously interacting components such as those implemented by services. From the formal point of view, agent-based logics, including knowledge and epistemic logics, have been used to reason about the composition of services [LQS08a]. Agent-based approaches have been used to monitor service executions. Of particular importance to the project are contributions made in the area of verification of multi-agent systems [GvdM04, PL03] and, especially their applications to the services domain. In this line of work model checkers for multi-agent systems, e.g., MCMAS [LQR09] have been used to verify web-services. For example in [LQS08a] a compiler is provided from WSBPEL into ISPL, a modeling language for agent-systems.

The starting point in verification in the artifact paradigm will be existing work on model checking [CGP99] and, specifically, three research directions: (i) Abstraction, (ii) model checking of multi-agent systems, and (iii) existing open-source implementation toolkits. At present there is only very preliminary work on verification of systems with artifacts [BGH+07, DHPV09]. Artifacts incorporate both data and processes, so an in-depth study will be required. In terms of abstraction, our work will build upon techniques for infinite state systems. Such techniques have been proposed in the past, e.g., data abstraction [AA99+, BLO08], abstract interpretation [DDP99, WB98], and partial evaluation [LM99]. In terms of model checking multi-agent systems, while current work includes bounded model checking via SAT [PL03] and Reduced Ordered Binary Decision Diagrams [RL07], these are based on systems with finite states so they will have to be suitably extended to deal with the artifact centric approach of interest here. Lastly, experimental model checkers for the verification of multi-agent systems environments have been developed [GvdM04, LQR09] and released as open-source. Given its support for ATL as well as epistemic logic, the model checker MCMAS [LQR09] will be considered as a starting point for the research to be carried out here.

Progress in ACSI beyond the state of the art: In ACSI, we will design, and evaluate verification techniques in the artifact paradigm. We will start by analysing suitable specification languages; we expect that plain temporal logics (both in their linear and branching time variants) will be insufficient to specify the kind of high-level properties we will need to use to model artifacts. Candidates for *extensions of specification languages* will be epistemic logics [FHMV95], deontic logics, ATL [AHK02], and possibly suitable combinations of them. Similarly, an investigation will be made as to whether the specification requirements will force us to use a quantified language, whether a propositional one suffices, or whether we need to adopt a position in between, e.g., by quantifying on propositions. The definition of the specification language will greatly influence further research. This involves devising *novel model checking algorithms* for artifacts. A crucial component will be the representation of the data structures to be used. Most likely we will use *efficient symbolic data structures* such as BDDs [Bry92] or more sophisticated variants, e.g., trees of BDDs. In addition to the above, we will face the task of discretising the state space resulting from the progression of artifact instances through their lifecycles. In fact, given that artifacts combine processes with data, the resulting state spaces will in general be infinite. To surpass this difficulty we will extend the previous work in model checking of infinite state systems mentioned above (data abstraction, abstract interpretation, and partial evaluation) in ways that are compatible with the specification languages devised. A guiding principle that we will follow here is that we will be aiming to achieve *implementable algorithms with adequate* *performance* in the artifact paradigm. We will also design and *implement a verification toolkit* implementing the model checking techniques for verification both in the artifact layer and in the semantic layer. The toolkit will build upon existing state-of-the-art model checkers (such as NuSMV [C+02] or MCMAS [LQR09]) by adapting and extending the verification algorithms to the case of artifact centric systems. We plan to release the toolkit as open source for maximum impact and potential reuse and extension by the community.

1.2.6 Process mining and monitoring

Existing service-oriented information systems record an abundance of events in the form of message logs, but these logs are not being fully exploited to maintain and evolve such systems [vdARW+07]. Process mining refers to a body of techniques that take as input such logs, and help to uncover information about the processes that these systems support [AGL98, vdAWM04,]. Examples of questions that can be answered by using today's process mining techniques include: *(i)* Conformance checking [CW99, vdADO+08,], i.e., checking whether regulations and contracts are followed by a running system, based on a stream of logs; *(ii)* Process discovery [vdAWM04,], i.e., reverse-engineering business process models from system logs; *(iii)* QoS monitoring, i.e., identifying performance bottlenecks, monitoring flow times, resource utilization, response times, error rates, etc. We envisage that conformance checking and process discovery will be two of the key pillars for monitoring and maintaining service interoperation hubs in the ACSI framework. A third pillar will be on-the-fly adapter generation, which will be used to fix sources of incompatibility detected by conformance checking techniques. Current approaches to adapter synthesis [MNBM+07] are geared towards offline adaptation, and cannot be readily applied to generate adapters while the system is running. Recently, some research on dynamic adapter generation for services has been undertaken [CSC08], but its results have not yet been thoroughly evaluated in large service networks.

Progress in ACSI beyond the state of the art: Current process mining techniques suffer from three major shortcomings that make them unsuitable in the context of artifact-centric service interoperation hubs.

First, existing process conformance checking techniques are designed to work post-mortem, meaning that the logs are collected, cleaned and analyzed off-line. In the context of service interoperation hubs with independently evolving services and dynamic relations, deviations with respect to the expected execution need to be detected and pinpointed in real-time. Accordingly, ACSI will develop *online conformance checking techniques that scale up* to hundreds or thousands of interactions per second, which are characteristic of large-scale interoperation hubs. Second, existing process mining techniques are focused on identifying control-flow dependencies. No comprehensive technique exists for mining data dependencies in processes. This means that existing process mining techniques could be adapted to reverse-engineer the lifecycle of a given artifact from a set of execution logs of the artifact. But these techniques cannot discover which services read and write which attributes of the artifacts and at which point in time. *Mining* such *read/write actions from service interaction logs*, will be one of the challenges to be addressed in ACSI. Finally, in the context of artifact-centric interoperation hubs, one has to monitor and discover the artifacts themselves and the interactions between these artifacts. Conformance checking and process discovery of such artifacts and interactions are beyond the scope of existing process mining techniques. In

ACSI, we will devise *techniques for discovering artifacts and interactions between artifacts from service interaction logs*.

On the adapter generation front, ACSI will advance the state-of-the-art by developing and evaluating techniques for adapter generation that, given the output of a conformance checking module, will efficiently generate an adapter to fix the detected incompatibilities. One of the key challenges here will be to *automatically establish relations between data elements in different services*. For this purpose, we will employ information contained in the semantic layer of artifact models.

1.2.7 Related projects

In the following table, we list both projects that have developed some of the technologies on which ACSI builds, and projects that are similar in scope to ACSI and for which ACSI adds a layer of innovation that goes well beyond what each of these projects promised to deliver.

Project's Name and Focus of Innovation	ACSI take-up and/or differentiation
<u>Services and Services Oriented Architecture</u>	
COMPAS <http://www.compas-ict.eu/> will design and implement novel models, languages, and an architectural framework including required software components and services to ensure dynamic and on-going compliance of software services to business regulations and the stated user service-requirements. This is achieved using the model-driven software development approach to enable organizations developing custom business compliance solutions faster, cheaper, and with less required programming skills.	*Similar to COMPAS, also in ACSI we can follow a model-driven approach, by compiling artifacts into running programs. However, artifacts abstract from most implementation details, which makes it possible to use this level of abstraction for formal verification of activities at design time, to check formal compliance to contracts at design and/or deploy time, and to adopt formally grounded observation-based techniques at runtime. The emphasis on the formal dimension is a key feature distinguishing ACSI from COMPAS.*
The objective of **SHAPE <**http://www.shape-project.eu/**>** is to support the development and realization of enterprise systems based on a Semantically-enabled Heterogeneous Service Architecture (SHA). SHA extends SOA with semantics and heterogeneous information system infrastructures (Web services, agents, Semantic Web Services, P2P, and grid) under a unified service-oriented approach building on a semantic middleware. To achieve this, SHAPE will develop a model-driven engineering tool-supported methodology.	*SHAPE focuses on providing a unified view to multiple information system platforms through a semantic middleware, but it does not focus on a principled approach for describing how different services use and share data. This makes it difficult to establish sustainable and meaningful interoperation among existing services, and challenging to maintain an interoperation environment in the face of dynamic business objectives and relationships. ACSI will develop an architecture based on artifacts, analogous to SOA but incorporating data explicitly, which will overcome these drawbacks.*
SOA4All <http://www.soa4all.eu/> will help to realize a world where billions of parties are exposing and consuming services via advanced Web technology. The outcome of the project will be a comprehensive framework and infrastructure that integrates four	*ACSI does not focus on Web technology for SOA, but similar to SOA4All it will rely on semantic based technology to abstract from any specific implementation technology. However, ACSI will go beyond the usage of semantic technology for service discovery, as envisioned in SOA4All, and*

Project's Name and Focus of Innovation	ACSI take-up and/or differentiation
complimentary and revolutionary technical advances into a coherent and domain independent service delivery platform: *(1)* Web principles and technology as infrastructure for service integration. *(2)* Web 2.0 to efficiently structure human-machine cooperation. *(3)* Semantic Web technology to abstract from syntax to semantics as required for meaningful service discovery. *(4)* Context management to process in a machine understandable way user needs that facilitates the customisation of existing services for user needs.	*will make the Semantic Layer, realized through the KAB, a key element of the artifact paradigm, holding information about both the static and the dynamic aspects of the system. Service inter-operation, one of the explicit objectives of ACSI, will be obtained by realizing a framework for interoperation hubs, supported by a Design Workbench and a Hub Engine.*
EzWeb <http://ezweb.morfeo-project.org/> will provide an open source reference implementation of standard technologies for the front-end web access layer in next-generation SOA and the future Internet of Services. The project proposes a new approach for making services available to end users that relies on the concept of a "web resource" as a basic building block. The users can search for these building blocks in a catalogue and combine them into a personal "workspace". These web resources can then be linked together by the user himself in order to create flows of information that are triggered by interacting with web resources.	*EzWeb concentrates on the front-end service access layer, and provides end users with advanced mechanisms and support for combining web resources, relying on existing technology for the back-end interoperation of services. Hence, the results and technology developed within ACSI, and in particular the fact that the focus is on both data and processes will complement those achieved within EzWeb. Moreover, ACSI will rely on the semantic layer to bridge the separation between the business and the IT levels, a notoriously problematic aspect in traditional SOA.*
Semantic Web Technologies and Ontologies	
MOST <http://www.most-project.eu/> will improve software engineering by leveraging ontology and reasoning technology. To reach this goal, a seamless integration of ontology technology into model-driven software development will be realized, resulting in ontology-driven software development. This concerns the integration of all involved elements (ontology and modeling languages, models, tools), as well as the development processes (process and reasoning guidance, traceability of models).	*The focus of MOST is model-driven and ontology-based software development in general, rather than the development of services and the interoperation of different services, as in ACSI. Moreover, ACSI differs from MOST since it relies on the artifact paradigm, in which specific attention is given to the modeling of data, and to the interaction of the services with the data.*
ReDSeeDS <http://www.redseeds.eu/> aims at creating an open framework consisting of a scenario-driven development method (precise specification language and process for the "how-to"), a repository for reuse and tool support throughout. The basic reuse approach will be case-based, where a reusable case is a	*ReDSeeDS is concerned with requirements-driven software development, and follows an approach based on reusable SW cases, which can be adapted and transformed to new requirements. Hence, both the goals of the project and the approach taken are different from those in ACSI, which aims at setting up a framework for service*

Project's Name and Focus of Innovation	ACSI take-up and/or differentiation
complete set of closely linked (through mappings or transformations) software development technical artifacts (models and code), leading from the initial user's needs to the resulting executable application. The project will combine and enhance state of the art in the areas of requirements engineering, meta-modeling, model transformation, and querying and inference techniques.	*interoperation and for the verification, synthesis, and mining of the involved services. ReDSeeDS is based on a modeling language including both requirements and transformations, while ACSI relies on artifacts complemented with the KAB as key elements for service interoperation.*
TONES <http://www.tonesproject.org/> aimed at studying and developing automated reasoning techniques both for ontology engineering tasks (i.e., design, maintenance, merging, and integration) and for operating with ontologies (i.e., run-time access to ontologies and interoperation of ontologies and data). As such, it has set the foundations for a widespread use of ontologies and of the associated inference techniques in tasks related to data management.	*The results achieved within TONES on ontology-based data management provide some of the formal foundations for the work in ACSI. Specifically, the Description Logics developed within TONES that allow for efficient query processing over ontologies, will provide the foundations for the data modeling language in the KAB. Indeed, such logics enjoy good computational properties not only w.r.t. query answering, but also w.r.t. instance level updates of the ontology, which is a key operation in the KAB.*
ONTORULE <http://ontorule-project.eu/> aims at developing and integrating all the required pieces of knowledge and technology to allow the acquisition of ontologies and rules from the most appropriate sources, including natural language documents, their separate management and maintenance, and their transparent operationalisation in IT applications. It will develop techniques and algorithms for the seamless integration of logic-based ontology languages with different kinds of rule formalisms, ranging from production rules to logic-based rule formalisms.	*The framework for the integration of ontologies and rules into hybrid knowledge bases proposed within ONTORULE will provide a further solid basis for the Knowledge and Action Base of ACSI. Rules can be seen as one of the mechanisms for implementing actions; hence, ACSI will be able to exploit the research on the semantic properties of hybrid knowledge bases and on the inference algorithms for them developed in ONTORULE.*
SUPER <http://www.ip-super.org/> has combined Semantic Web Services and Business Process Management, into one consolidated technology. Specifically, it has created horizontal ontologies, which describe business processes; vertical tele-communications oriented ontologies to support domain-specific annotation for the telecommunication sector; and a suite of tools that bridge the business and IT domains.	*ACSI shares with SUPER the use of semantics to bridge the gap between the business and the IT domain. However, modeling, execution, and analysis of business processes in SUPER is carried out without focussing specifically on data management, which instead is one of the innovative features of ACSI. ACSI will be able to exploit ontologies such as those developed in SUPER as a basis for the semantic layer.*

Project's Name and Focus of Innovation	ACSI take-up and/or differentiation
Interoperation and distributed systems	
ATHENA <http://www.eic-community.org/> has been a comprehensive research initiative in IT to remove barriers to interoperability, to transfer and apply the research results in industrial sectors, and to foster a new networked business culture. ATHENA aimed to enable interoperability by providing a comprehensive Interoperability Framework. Research and Development have been executed in synergy and in collaboration with Community Building: research was guided by business requirements defined by a broad range of industrial sectors and integrated into piloting and training. ATHENA has been a source of technical inventions for interoperability, and led to prototypes, technical specifications, guidelines and best practices that form a common European repository of knowledge.	*ACSI will build upon the experience gained in ATHENA and the standards for business interoperation developed in the project. ACSI has in common with ATHENA the notion of a holistic framework for interoperability among enterprises. However, while ATHENA adopts a three layer framework for interoperation (the business layer, the knowledge layer, and the ICT layer) in a SOA environment with service orchestration, ACSI introduces two innovative elements to cope with service interoperation among enterprises: interoperation hubs and (dynamic) artifacts. The ACSI technology enables service collaborations without the traditional need of layer separation between business and ICT in an open service network.*
CONTRACT <http://www.ist-contract.org/> developed frameworks, components and tools to model, build, verify and monitor distributed electronic business systems on the basis of dynamically generated, cross-organisational contracts. These electronic contracts underpin formal descriptions of the expected behaviours of individual services and the system as a whole. They are verified offline (i.e., at design time) by using basic model checking technology.	*ACSI will build on and extend the model checking technology developed and used in CONTRACT, by incorporating also data in the form of artifacts in the formal specification, and by including the semantic layer constituted by the KAB. Both of this requires to substantially extend the techniques used in CONTRACT to deal with the infinite state space resulting from the presence of data.*
ADMIRE <http://www.admire-project.eu/> will accelerate access to and increase the benefits that can be gained from data exploitation by delivering consistent and easy to use technology for extracting information and knowledge from multiple heterogeneous and distributed resources. To cope with complexity, change, and heterogeneity of services, data, and processes, an abstract view of data mining and integration will be provided, giving power to users and developers of data mining and integration processes.	*ADMIRE concentrates on the technology for data mining and integration, by developing an infrastructure of gateways connected together over the Internet and Grid, each providing a core set of data mining and integration services, which can be driven using a high-level language. However, it does not address the general problem of service interoperation taking into account explicitly both the process and the data manipulated by the services, which is the focus of ACSI.*

1.2.8 Indicators

The ACSI project addresses target outcome (b) of Objective 1.2: Highly Innovative Service/Software Engineering. Henceforth we detail the list of indicators related to this objective to be used during the three years of the project. Collaboration target indicators will be filled out after the first collaboration meeting.

Project level		**Collaboration activities**	
Target	**Achieved**	**Target**	**Achieved**
Contribution to Standards -1 membership to a standard body (SBVR)			
Peer reviewed articles - 2 in one of the following journals (*)			
Panel discussions and key notes in conferences or workshops – 1 in one of the following conferences (**) or associated workshops			
Conference publications – 2 in two of the following conferences (**)			
Dissemination towards general public – 2 press releases			
Tutorials at conferences or workshops – 1 tutorial based on ACSI Technologies in one of the following conferences (**) or associated workshops			
Academic courses – 1 graduate course based on ACSI Technologies			

(*) IEEE Transactions on Services Computing, Elsevier journal Computer for Industry, IEEE Intelligent Systems, ACM Transactions on Database Systems (TODS), ACM Transactions on Computational Logic (TOCL), Journal of Computer and System Sciences (JCSS), Information and Computation (IC), Artificial Intelligence Journal (AIJ), Journal of Artificial Intelligence Research (JAIR), Information Systems (IS), Software and System Modeling, Journal of Cooperative Information Systems, ACM Transactions on Internet Technology.

(**) Conference of Association for the Advancement of Artificial Intelligence (AAAI), International Conference on Knowledge Representation and Reasoning (KR), International Joint Conference on Artificial Intelligence (IJCAI), International Conference on Principles of Database Systems (PODS), International Conference on Database Theory (ICDT), International Conference on Very Large Data Bases (VLDB), International Conference on Service Oriented Computing (ICSOC), Business Process Management Conference (BPM), International Conference on Autonomous Agents and Multiagent Systems (AAMAS), International Conference on Automated Planning and Scheduling (ICAPS), IEEE International Conference on Web Services (ICWS), IEEE International Conference on Services Computing (SCC), IFIP International Conference on Formal Techniques for Distributed Systems (FMOODS/FORTE), OnTheMove Federated Conferences & Workshops (OTM), Computer Aided Verification (CAV),), International Semantic Web Conference (ISWC).

The following table describes Objective 1.2 (b) impacts and ACSI's foreseen impact indicators. Collaboration indicators should be described at a later phase of the project.

Project level		**Collaboration activities**	
Target	**Achieved**	**Target**	**Achieved**
Contribution to Future Internet / Convergence			
The ACSI interoperation hub will provide a new way in the conceptualization of future service collaborations and value nets.			
Technological advances in software/service engineering			
A new framework for service interoperation as an alternative to the now popular orchestration and choreography			
More competitive environment for service providers, including SMEs			
The ACSI interoperation hub framework will enable a flexible scalable collaboration environment.			
Massive uptake of high- added value services. Service Front-ends, online communities			
The ACSI interoperation hub will make it faster and less expensive to create new value networks that can bring together collections of services (and service front ends)			
Strengthened European software and services industry			
Much of the software prototyping for this project will be performed in an open source manner. The ACSI interoperation hub can be applied to essentially any domain.			

B1.3 S/T Methodology and associated work plan

1.3.1 Overall project strategy

Project strategy distinguishes between scientific and technology research activities, which are carried out in parallel. Scientific research is focused on the development of a new framework, techniques, and tools for enabling service interoperation based on interoperation hubs, and research on KABs, verification, and process mining. The technology research will take place over three one-year phases. In the first phase, the interoperation hub architecture will be established, prototyped, and tested extending an existing (state-machine based) approach to artifacts. The second phase will incorporate principled methods for designing interoperation hubs, incorporating the outputs of the scientific research available at that time. The final stage will incorporate the remainder of the scientific research output. Each phase will be tested in the context of two real-world domain applications regarding Energy Distribution Management and European Journal. Both the scientific and technology research will be guided by the requirements revealed by these two scenarios. ACSI goals and specific research and technical innovations will be carried out using a compact structure made of six well coordinated RTD work packages and a management work package. The work packages will interrelate as described in Figure 1-7. The work packages cover all aspects of the three project streams, namely the scientific, technology, and validation streams (see Figure 1-8). Table 1-1 below provides an overview of the proposed risk management and contingency plans.

WP1, ***ACSI artifact paradigm and interoperation hub framework*** **(led by Uniroma1)** will focus on development of the interoperation hub framework, based on the end-to-end artifact-centric approach. The framework will include the new semantic layer for artifacts, and "hooks" where the results of the other scientific research work can be incorporated.

WP2, ***Formal-based techniques and tools*** **(led by Imperial)** will concentrate on formal methods and tools for supporting design and realization phases of artifact-centric interoperation hubs. A theory of Knowledge and Action Bases (KABs) will be developed, to facilitate the modeling of the semantics of artifact-centric services together with their possible evolution.

WP3, ***Observation-based techniques and tools*** **(led by UT)** will focus on automated extraction and verification of services, artifact types, and interoperation hubs from service execution traces, both in online and offline settings. Also provided in this work package are the means for conformance checking of services against artifact models, and on-the-fly adapter synthesis of adapters for resolving mismatches.

WP4, ***Development of ACSI Hub System*** **(led by IBM)** will be responsible for the development of the ACSI Hub System (mainly the Design Workbench and the Interoperation Hub Engine) along with the integration of the modules developed in WP2 and WP3.

WP5, ***Validation*** **(led by INDRA)** The ACSI Hub System will be validated in the context of two real-world domain applications. The first use case is focused on managing electrical power distribution, while the second is focused on a European Journal.

WP6, ***Dissemination, exploitation and communication (led*** **by IBM)** will coordinate all the dissemination and exploitation activities related to the project.

WP7, ***Project management*** **(led by IBM)** will coordinate all the management activities related to the project.

Interrelationships among work packages (PERT diagram)

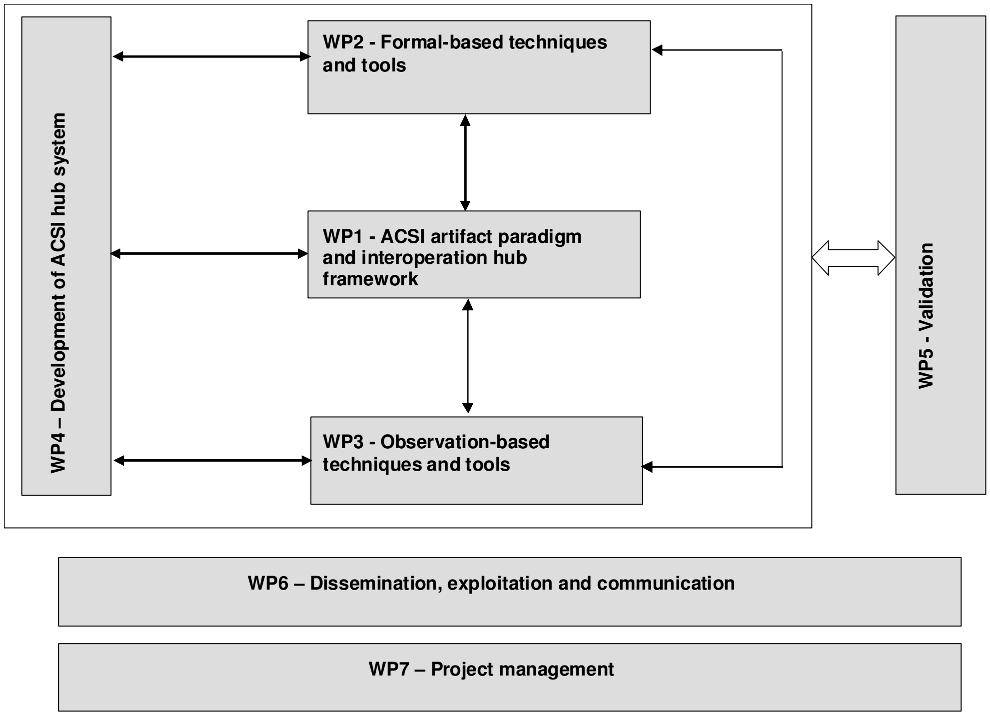

Figure 1-7: Pert diagram

Mapping of project streams into work packages

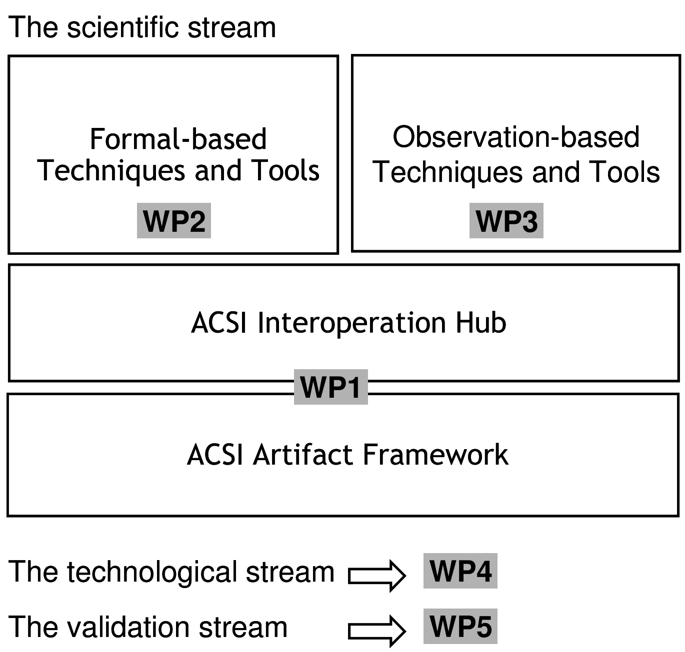

Figure 1-8: Mapping of project streams into work packages

Table 1-1: Project risk management and contingency plans

#	WP #	Risk	Probability (Medium, Low, High)	Impact (Medium, Low, High)	Risk management / contingency plan
1	**7**	One partner fails	**Low**	**Medium**	Carefully selected partners, with previous cooperation and experience in EU projects (see section 2.3).
2	**7**	Failure to coordinate and follow up on project progress and work plan, resulting in delays and failed tasks.	**Low**	**High**	The prime contractor has significant experience in managing international projects, including the coordination of EU projects (e.g. Reservoir, ModelPlex, PROSYD, SHADOWS, SAPIR). A strong management structure has been established including control measures to assure progress according to the work plan.
3	**1,2, 3,4**	Some of research tasks appear too complex	**High**	**Low**	Working through iteration process to deal with complexity.
4	**1,2, 3,4**	Lack of consensus on the technological approach between competence areas	**Low**	**Medium**	WP4 acts as an integration WP, providing common technological infrastructure for other WPs. WP4 also serves as an integration point bringing together the results of the other WPs. Thus, WP4 will provide the appropriate setting for ensuring uniformity of the technology approach employed across the project.
5	**1,2, 3,4, 5**	Quality of technical results is not good enough	**Low**	**Medium**	The WPs proceeds incrementally, in cycles of development and subsequent refinement, with checkpoints at different milestones. Should the initial results be technically insufficient, this will be detected early in the project so that appropriate measures can be taken.
6	**5**	The project results do not achieve sufficient industrial relevance	**Low**	**Medium**	The results of the project will be evaluated in two different industrial case studies in two different domains, and led by different industry partners. This will provide different channels for exploring potential applications of the produced technology and for assessing the industrial relevance of

#	WP #	Risk	Probability (Medium, Low, High)	Impact (Medium, Low, High)	Risk management / contingency plan
					the results.
7	**6**	The dissemination of the project results is not sufficient to create impact	**Low**	**Medium**	A significant amount of resources are allocated in the work-plan to dissemination. Furthermore, the project partners have significant experience in dissemination both via academic and industry channels. The likelihood of this event is thus low.
8	**1**	Tasks to carry out on artifact defined in WP1 are not sufficiently understood by all participants, thus negatively influencing activities of WP2, WP3 and WP4	**Low**	**High**	High involvement of all participants in tasks' definitions.
9	**1**	Failure in identification of problems related to the semantic-layer that admit practical implementable solutions	**Low**	**High**	This point will be investigated since the very beginning of the WP, and a great effort and attention will be posed on it.
10	**1**	The conceptual architecture of the interoperation hub does not perfectly reflect the ACSI paradigm and the tasks devised are not aligned with artifact tasks	**Low**	**High**	The definition of the interoperation hub framework will be based on and strongly influenced by the definition of the artifact layer and realization layer.
11	**5**	Difficulties on the analysis of requirements of the actual system and delay related problems coming from other package developments	**Medium**	**High**	A set of different alternatives have been established if a delay from other packages occur. We plan to have monthly meetings with each use case team to update the state of each pilot and to have regular updates to detect possible deviations.
12	**4**	Deliverables from WP1, WP2, and WP3	**Medium**	**Medium**	Strong management structure to ensure deliverables are carried out in

#	WP #	Risk	Probability (Medium, Low, High)	Impact (Medium, Low, High)	Risk management / contingency plan
		and not received in time, or are not as expected			time. Strong interrelations among the WPS and the involvement of many of the partners ensure deliverables are carried out and delivered as expected. Moreover, deliverables of crucial components will have several versions.
13	**4**	ACSI Hub System fails	**Medium**	**High**	The project will adopt a three phase iteration until the final prototype system. Each iteration will include lessons learnt from previous iteration tool and its validation on the use cases. Furthermore, first iteration includes extension to the state-of-the art proven technology.
14	**6**	Failure to widely expose the project results to researchers and stakeholders	**Low**	**High**	A strong management consortium that comprises world leaders in their expertise and good support relationships between research partners and industrial partners to guarantee the exposure of the project to potential stakeholders.

1.3.2 Timing of work packages and their components

Work packages, Tasks, Deliverables & Milestones	
WP1	**ACSI artifact paradigm and interoperation hub framework**
T1.1	Definition of the artifact-layer and realization-layer
T1.2	Definition of the semantic-layer
T1.3	Definition of the interoperation hub framework
WP2	**Formal-based techniques and tools**
T2.1	Artifact verification and synthesis techniques
T2.2	Model checker for artifact interoperations
T2.3	Techniques and tools for KAB, to manage data linkage with artifact layer
T2.4	Techniques and tools for KAB, to manage action linkage with artifact layer
T2.5	Evolvability at the interoperation hub
WP3	**Observation-based techniques and tools**
T3.1	Conformance checking of services against artifact models
T3.2	On-the-fly adapter synthesis
T3.3	Discovery of artifact lifecycles
T3.4	Discovery of data access and update rules in artifact-centric models
T3.5	Discovery of artifacts in interoperation hubs
WP4	**Development of ACSI hub system**
T4.1	ACSI Hub System Iteration I
T4.2	ACSI Hub System Iteration II
T4.3	ACSI Hub System Iteration III
WP5	**Validation**
T5.1	Energy Distribution Management and European Journal use cases definition and requirements
T5.2	AS-IS assessment
T5.3	Pilots based on the release of ACSI Hub System iteration I
T5.4	Pilots based on the release of ACSI Hub System iteration II
T5.5	Pilots based on the release of final ACSI Hub System
WP6	**Dissemination and exploitation**
T6.1	Dissemination action plan
T6.2	Collaboration with ICT IoS Projects
T6.3	Publications
T6.4	Events
T6.5	Future development
T6.6	Press releases
T6.7	Project story - factsheet
T6.8	Public website
T6.9	Communication materials
WP7	**Project management**
T7.1	Project organization and planning
T7.2	Technical coordination
T7.3	Quality of control
T7.4	Reporting
T7.5	Administrative and financial coordination
T7.6	Internal project Web site
T7.7	Coordination of dissemination activities

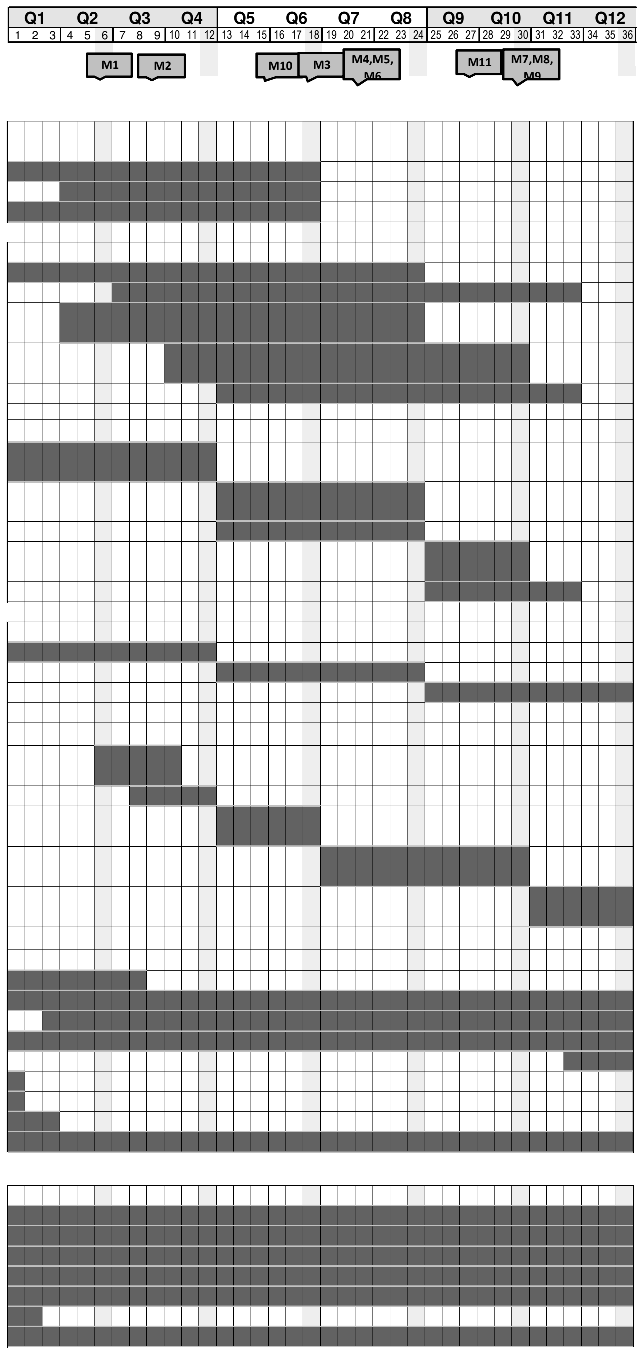

Figure 1-9: Gantt chart

B2. IMPLEMENTATION

B 2.1 Management structure and procedures

2.1.1 Project management

The aim of the project management is to guarantee that the objectives of the project are achieved on time, on budget, and with high quality. The ACSI project will be managed with sound and efficient decision making, execution, and control—and will maximize partner accountability, commitment, involvement, and prospects of success.

To implement the above goals, the proposed project management structure includes the following figures:

- WP Leader
- Technical Steering Committee
- WP Team

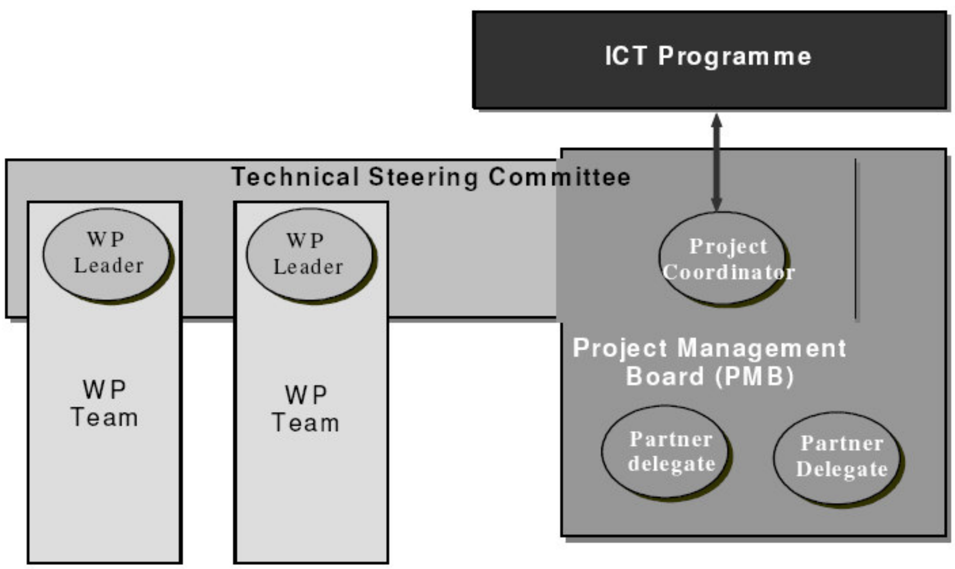

Figure 2-1 Project management structure

The **project coordinator** (Dr. Fabiana Fournier, from IBM Haifa Research Lab), acting as the representative of the prime contractor, will be the permanent reference point of the ACSI project on a day-to-day basis vis-à-vis both the contractors and the European Commission (EC). The project coordinator is in charge of the overall project organization, planning, and reporting. The project coordinator shall specifically have the following functions:

- reporting to the commission and serving as the administrative liaison with the EC;
- monitoring project progress and workload consumption;
- chairing the Project Management Board (PMB) meetings, and following up on implementation of PMB decisions;
- presenting the overall project and the cost statement to the audit meetings.

- monitoring the Parties' compliance with their obligations under the Grant Agreement and the Consortium agreement.
- reviewing the reports to the EC to verify consistency with the Project tasks before transmitting them to the EC. Every deliverable will be submitted to an internal review by at least two consortium members.

The Project Management Board (PMB) will consist of official delegates assigned by the project partners, and will be chaired by the project coordinator. The management representatives will have the authority to make decisions on behalf of their respective organizations in terms of overall strategy and resources allocated to the project. The PMB will determine the strategic direction of the project and will be in charge of the high level management of the project, addressing all the administrative, contractual, and financial matters. It will vote on all important decisions related to the contractual execution such as changes to the consortium configuration, reallocation of work, responsibilities and man-power between contractors, settlement of problems, or differences between contractors.

Table 2-1: Project Management Board (PMB)

Partner	Delegate
Project coordinator	Dr. Fabiana Fournier (chair)
Uniroma1	Prof. Giuseppe De Giacomo
UT	Prof. Marlon Dumas
TU/e	Prof.dr.ir. Wil van der Aalst
FUB	Prof. Diego Calvanese
Imperial	Dr. Alessio Lomuscio
Indra	Daniel Martinez Maqueda
Collibra	Dr. Pieter De Leenheer

From a technical point of view, the project is broken down into a number of **work packages** (WP), each of them addressing a specific area of work. The practical work of a WP team will be conducted by contributions to periodic meetings and by partial contributions which will be submitted by individual members for formal adoption by the group. Each WP will be coordinated by a **WP leader**. Each WP is further subdivided into its large component tasks, which are allocated a **Task Leader** responsible for coordination.

The technical coordination among WPs will be handled by the **Technical Steering Committee** (TSC), which will be composed of all WP leaders and chaired by the project coordinator.

In particular the TSC will be responsible for the implementation of the directives of the PMB, the guidance and monitoring of the technical WPs, and the coordination among WPs, the timely preparation, approval, and forwarding to the Commission of the deliverables produced by the WPs, and the resolution of conflicts amongst WPs.

Table 2-2: Technical Steering Committee

Work package	Work package Leader/Representation
Technical Coordinator	Dr. Fabiana Fournier / ACSI project coordinator
WP1 ACSI Artifact and Interoperation Hub Framework	Prof. Giuseppe De Giacomo/Uniroma1
WP2 Formal-based techniques and tools	Dr. Alessio Lomuscio / Imperial
WP3 Observation-based techniques and tools	Prof. Marlon Dumas / UT
WP4 Development of ACSI Hub System	Dr. Fabiana Fournier / IBM
WP5 Validation	Daniel Martinez Maqueda / Indra
WP6 Dissemination and exploitation	Dr. Fabiana Fournier / IBM
WP7 Project Management	Dr. Fabiana Fournier / IBM

If necessary, the PMB can create ad-hoc Task Forces, composed of experts, chosen from the project participants, that will work together to solve well-defined problems in a limited period of time.

2.1.2 Meetings and decision process

The PMB and TSC will meet every three months to check and supervise the progress of the project. In the initial phases of the project, meetings will include quick-start workshops to produce detailed technical architectures, finalise work plans and establish partner responsibilities. Detailed technical and organisational work will be done between meetings as much as possible, using the available communications infrastructure (e.g. e-mail, a project WWW site, and group audio and videoconferencing facilities).

Decisions of the PMB and TSC are made by a simple majority of members of the consortium (when not explicitly stated otherwise). Each member has one vote, which may be cast via a proxy if they are unable to attend the meeting.

As for active WPs, periodic meetings will be organized depending on the work to be carried out. Two kinds of WP meetings are foreseen:

1. Planning meetings, where all relevant decisions about the work to be carried out within the WP will be taken.
2. Integration meetings, where WP members will sit down in front of terminals to integrate and finalise their contributions

2.1.3 Consortium agreement

The relationship between all contractors will be fixed in a ***Consortium Agreement*** based on the following principles:

In order to have a management system applicable through all phases of the project, a reasonable approach is to have straight, clear and direct management and organization protocol at all levels. This is particularly relevant given the challenging financial and industrial policy constraints. Therefore, in order to have clearly assigned responsibilities, to avoid any friction and to progress as per the project plan, the responsibilities and authorities of the project manager and the team members will be unambiguous.

The partners will be bound by a formal consortium agreement that is planed to be signed prior to the beginning of the project of the project, and in which their roles, responsibilities and mutual obligations will be defined both for the project life.

2.1.4 Conflict resolution

When conflicts arise during the execution of the project, they will be resolved according to the following principles:

- They will first be addressed within the relevant WP through discussion chaired by the WP leader.
- If this fails, the issue will be presented by the WP leader either to the technical steering committee or the project management board depending on the nature of the problem (technical or business/strategic).
- The relevant board will attempt to resolve the issue through a majority vote.

Technical issues between WPs will also be addressed by the technical steering committee. As noted above, the technical steering committee of the project consists of the WP leaders (chaired by the technical coordinator), and the project management board consists of representatives of each partner (chaired by the management coordinator). Any conflicts that cannot be resolved through the principles above will be handled according to the dispute resolution provision set forth in the Consortium Agreement.

2.1.5 Risk management

As a scientifically technologically challenging initiative, the ACSI project carries a degree of risk, which the partners aim to control. During the quarterly meetings, the PMB will hold a dedicated session to identify, evaluate, and track project risks.

The risk management strategy of the project will involve the following aspects:

- Risks, generically, will be independently analyzed across the three risk dimensions: cost, time, and quality of results.
- We will provide initial results frequently to early feedback groups in order to identify possible issues with the deliverables.
- The partners, under the responsibility of the WP leaders, will strive to increase transparency in regard to dependencies and exposures (activities, time, resources, and cost).
- Contingency plans and/or plans for corrective actions will be developed and implemented, should a risk be identified.
- Each work package has a dedicated risk assessment/management strategy that targets the risk associated with the specific work package. These WP-level risk management strategies are detailed in section 1.3.6.

2.1.6 Raising public awareness

Public awareness will be addressed in two major ways. Firstly, the project will have a web site for posting project activity and achievements. The site will be accessible to the public. Web site construction and management are part of WP6, Task 6.2. Secondly, the members of the project will publish technical and scientific results at scientific events. These events are open to the public and attract the most relevant audience, thus increasing public awareness exactly where this awareness is necessary. The papers to be published are similar in content and structure to research reports that are to be delivered as part of ACSI work packages. Deliverables from Work package 6 will include reports on these publications. ACSI will also take advantage of Web 2.0 techniques to increase awareness, e.g., twittering of news feeds when tool updates are put online or publications are accepted at conferences; blogging intermediate results and announcements; discussion fora; social networking environments (LinkedIn, Twitter); etc. As part of these events, the project will organize an annual workshop for both scientific and industrial communities.

B 2.2 Beneficiaries

2.2.1 IBM Israel – Science and Technology LTD (IBM, Coordinator)

IBM has the world's largest IT research organization, with more than 3,000 scientists and engineers working at eight labs in six countries. IBM leads the world in patent filing, surpassing the combined total awarded to 12 of the largest IT companies. In aggregate, the company holds nearly 37,000 patents worldwide. IBM Israel Science and Technology Limited is better known as the IBM Research Laboratory in Haifa (HRL). Since it first opened as the IBM Scientific Centre in 1972, HRL has conducted decades of research that have been vital to IBM's success. R&D projects are being executed today in areas such as cloud computing, smarter telco, healthcare and life sciences, discovery, verification technologies, multimedia, active management, information retrieval, programming environments, business transformation, and optimization technologies. The Lab houses IBM's biggest research centre outside the US, and employs over 500 people. HRL has strong legacy in leading and contributing to EU projects such as ModelWare and ModelPlex.

Role attributed in ACSI

IBM is the coordinator partner for the ACSI project. IBM will provide both management and technical contributions to the project. IBM will lead the development of ACSI Hub System (WP4) and deeply contribute to WP1 and WP2. IBM will also take a major role in validation (WP5).

Key Personnel

Dr. Fabiana Fournier has a B.Sc. in Information Systems, and a M.Sc. and a PhD in Industrial Engineering from the Technion - Israel Institute of Technology. Dr. Fournier has over ten years of research and practical experience in the areas of organizational and process modeling, business transformation, and management. In 2006, Fabiana joined the Business Transformation group at the IBM Research Lab in Haifa, Israel. Her research areas include: architecture and modeling (conceptual and formal modeling), model management, modeling tools architecture, enterprise architecture, and service oriented architecture**.** In 2008, Dr. Fournier received the IBM Outstanding Accomplishment Award for her contribution to the area of Business Architecture in the CBM (Component Business Modeling) domain in which she has published multiple articles in leading conferences and journals.

Dr. Richard Hull is a Research Staff Member and Manager at the IBM T.J. Watson Research Center, a position he took in May, 2008. Dr. Hull has broad research interests in the areas of data and information management, workflow and business processes, and web and converged services. His group is currently working performing research on data-centric workflow and business process management. A primary focus of his research is on the use of a novel marriage of data and process, called "business artifact", as a unifying foundation for the management of business operations. Dr. Hull is co-author of the book "Foundations of Databases" (Addison-Wesley, 1996); has published over 100 articles in journals, conferences and books; and holds six U.S. patents. Prior to joining IBM Research, Dr. Hull was Director of Computing and Software Principles Research at Bell Labs Research, a division of Alcatel-Lucent. While there, in addition to pursuing research on semantic web services, converged services, personalization, and data management, Hull was instrumental in developing and transferring new technologies into Alcatel-Lucent's product line, including the Vortex policy engine and the Datagrid data integration tool. Before joining Bell Labs in 1996 he served on the faculty of Computer Science at the University of Southern California, and was a frequent visitor at INRIA in France. His research has been supported in part by grants from NSF, DARPA, and AT&T. Dr. Hull was named Bell Labs Fellow in 2005 and ACM Fellow in 2007.

2.2.2 Università degli Studi di Roma La Sapienza (Uniroma1)

The Department of Computer and System Sciences (DIS) at Università degli studi di Roma "La Sapienza" is a center for research and education at the undergraduate and graduate levels. In computer science and engineering DIS is one of the world leaders in academic research on Artificial Intelligence, Knowledge Representation, Databases, and Algorithms. The Database and Knowledge Representation group, which is involved in the present project, is one of the best known research groups in Description Logics (DLs), data integration, service/process composition and integration, and reasoning about actions. In the above mentioned areas, the group achieved very relevant results, in particular concerning the classification and characterization of architectures for information integration and techniques for query processing in this context, as well as, service/process modeling and automatic service composition, or conceptual data modeling. Furthermore, some of its members are considered among the founders of the research in the area of DLs, which is currently considered of crucial importance for the development of semantic-oriented technologies (Semantic Web). Among its members it includes the following faculties: Prof. De Giacomo, principal investigator of the Uniroma1 unit within the present project, Prof. Maurizio Lenzerini, Prof. Tiziana Catarci, Prof. Riccardo Rosati, Prof. Giuseppe Santucci, Ing. Massimo Mecella, and Ing. Domenico Lembo. Group members serve on a regular basis as program committee (PC) members of the most prestigious conferences in the area of Artificial Intelligence and Databases (e.g., PODS, ICDT, IJCAI, KR) and have been often invited speakers at various international conferences, or PC or general chairs of numerous events including, e.g., ACM-PODS-08.

Role attributed in ACSI

The main contributions expected from Uniroma1 are: foundational research on the main issues of the project; support for all representational and reasoning aspects of the knowledge and action base, with a specific focus on the action component; verification and synthesis of processes that deal with data. Uniroma1 will lead WP1 and strongly contribute to WP2.

Key personnel

Prof. Dr. Giuseppe De Giacomo (www.dis.uniroma1.it/~degiacomo) is full professor in Computer Science and Computer Engineering. His research interests include service composition, information integration, knowledge representation and reasoning, reasoning about actions, cognitive robotics, and object-oriented methodologies. He has been the principal investigator for the unit of "La Sapienza" of the EU Project EU-FP6-FET TONES, and he is the scientific coordinator of the Open Collaboration – Open Source Agreement W0954341 between "La Sapienza" and IBM T. J. Watson Research Center. He is the author of more than 200 publications in international conferences and journals, including the most prestigious ones in the above mentioned areas, such as Journal of Computer and System Science, Information and Computation, Artificial Intelligence, Information Systems, ACM Transactions on Computational Logic, Theoretical Computer Science, PODS, SIGMOD, LICS, ICDE, VLDB, ICDT, IJCAI, AAAI, KR, ECAI, ICSOC, BPM. His h-index is 49 (one of the highest in Italy in the area of computer science and engineering) according to Google Scholar. He is associate editor of the Journal of Artificial Intelligence Research (JAIR).

Prof. Dr. Maurizio Lenzerini (http://www.dis.uniroma1.it/~lenzerini) is full professor in Computer Science and Computer Engineering. He is the author of several academic books on fundamentals of Computer Science, Software Engineering, and Database design. His main research interests include conceptual data modeling, data integration, data warehousing, semistructured data management, knowledge representation and reasoning, and object-oriented methodologies. He is the author of more than 300 publications in international conferences and journals. His h-index is 53 according to Google Scholar (the highest in Italy in the area of computer science and engineering). He is member of the Editorial Board of various international journals, in particular is area editor of Information Systems. He organized several international conferences and workshops and he was Program co-chair of CoopIS'99, Conference Chair of ER'99, and Program Chair of ICDT'03 and PODS'08.

2.2.3 Libera Università di Bolzano (FUB)

The Faculty of Computer Science at the Free University of Bozen-Bolzano began its teaching and research activities in 2001. The Faculty is a research intensive environment where both basic and applied research is carried out in collaboration with the international research community and the local environment. The KRDB Research Centre for Knowledge and Data at the Faculty of Computer Science was founded in 2002, and it now comprises about 25 researchers, including PhD students. The center aims at being an international center of excellence in basic and applied research on knowledge representation and database technologies and at proposing to selected enterprises innovative ideas and technologies based on the developed research. The research topics include Conceptual Data Modeling and Ontology Design, Intelligent Information Access and Query processing, Information Integration, Peer to Peer systems, Semi-structured Data, Distributed and Web Information Systems, Web Services, Computational Logic, and Logic-based Computational Linguistics. The KRDB Research Centre has been involved in many national and European projects, including the EU funded projects Esprit 22469 DWQ, IST-2001-34825 SEWASIE (www.sewasie.org), FP6-507482 KnowledgeWeb (knowledgeweb.semanticweb.org), FP6-506779 REWERSE (www.rewerse.net), IST-508011 INTEROP, FP6-7603 TONES, and FP7-231875 ONTORULE (ontorule-project.eu). The centre organises also the European Erasmus Mundus Master in Computational Logic and the European Erasmus Mundus Master in Language and Communication Technologies. The members of the KRDB Research Centre mostly involved in ACSI will be Diego Calvanese (PI), Alessandro Artale, Werner Nutt, Enrico Franconi, Mariano Rodriguez-Muro, and several postdocs and PhD students.

Role attributed in ACSI

FUB will play an important role in WP1 "ACSI artifact paradigm and interoperation hub framework", specifically in the definition of the artifact-layer and of the semantic layer. It will also play a prominent role in WP2 "Formal-based techniques and tools", specifically for techniques and tools for data linkage between the Knowledge and Action Base and the artifact layer, and for evolvability of the interoperation hub.

Key personnel

Prof. Dr. Diego Calvanese is associate professor at the Faculty of Computer Science, Free University of Bozen-Bolzano, Italy, where he teaches graduate and undergraduate courses on knowledge bases and databases, ontologies, knowledge representation, theory of computing, formal languages, and information integration. He is vice-director of the KRDB Research Centre for Knowledge and Data, and his research interests include formalisms for knowledge representation and reasoning, ontology languages, Description Logics, conceptual data modeling, data integration, web service modeling and composition, and semistructured and XML data management. He has been principal investigator or co-investigator in several national and international research projects, and he has been the coordinator of the EU STREP FET Project "Thinking ONtologiES" (TONES). He has been co-chair of the Int. Workshops on Description Logics DL-2003 and DL-2007, and PC co-chair of the Int. Conference on Web Reasoning and Rule Systems RR-2008. He is the author of more than 150 refereed publications, including ones in the most prestigious international journals and conferences in the areas of databases and artificial intelligence, such as Journal of Computer and System Science, Artificial Intelligence, Information Systems, ACM Transactions on Computational Logic, Theoretical Computer Science, PODS, SIGMOD, LICS, ICDE, VLDB, ICDT, IJCAI, AAAI, KR, ECAI, ICSOC. He is one of the editors of the "Description Logics Handbook". His h-index is 47 according to Google Scholar.

Dr. Alessandro Artale is an Assistant Professor at the KRDB Research Centre for Knowledge and Data of the Faculty of Computer Science, Free University of Bozen-Bolzano, Italy. He got a PhD in Computer Science from the University of Florence in 1994. He published more than 60 papers in international journals and conferences. He acted as both Chair and PC member in conferences, and as editor of both proceedings and journal's special issues. He has 16th Int. Symposium on Temporal Representation and Reasoning (TIME-2009). His research has been funded by the European Community and by National funds. His main research subject concerns description logics, temporal logic, automated reasoning, ontologies and conceptual modeling. A particular emphasis is devoted to the formalization of conceptual modeling tasks in domains with high semantic complexity and characterized by a dynamic aspect.

2.2.4 Imperial College of Science, Technology and Medicine (Imperial)

The department of Computing at Imperial College London employs over 50 academic staff, about 55 postdoctoral researchers and 125 doctoral students. In the UK Government rating system the department is rated excellent in teaching, was ranked 2nd in the UK-wide Research Assessment Exercise of 2008, and received the top-most score in the previous Research Assessment Execercises (1996, 2001). Its main research areas are logic-based foundations of computing, software engineering, artificial intelligence, high-performance computing, and image processing. Specifically on the intended contribution to the project, the Department of Computing has been a leading international centre for research in artificial intelligence and computational logic since 1975, and in particular in knowledge representation and automated reasoning, including recent results in automatic model checking. The project will build directly upon this latter expertise.

At Imperial the project will be hosted by the Logic and Artificial Intelligence Section. This comprises 10 permanent members of staff and several post-doctoral and doctoral researchers. Throughout the years, the section has been supported by a number of research grants from the UK, the EU, as well as information technology companies.

Role attributed in ACSI

Imperial will offer key competences to the project in formal logic for multi-agent systems, verification of distributed systems and multi-agent systems. Imperial will lead WP2 (Formal-based techniques and tools). Imperial's contribution will focus on the definition and implementation of abstraction-based techniques for artifact centric environments. Imperial will follow actively a number of tasks in WP1 and also be active in WP4 contributing to the integration of the techniques developed.

Key personnel

Dr. Alessio Lomuscio is Reader in logics for multi-agent systems at the Department of Computing, Imperial College London. He received a PhD in Computer Science from the University of Birmingham in 1999 and a Laurea in Electronic Engineering from Politecnico di Milano in 1995. His research interests concern the specification and verification of distributed systems and multi-agent systems by means of techniques built on computational logic. He has published about 80 research papers on the subject. He was invited speaker and invited guest lecturer in a number of international conferences and postgraduate courses and regularly sits on programme committees for several international conferences and workshops. He is currently serving as principal Investigator for the EU Marie Curie FP7 project "First-order Modal Logics for the Specification and Verification of Multi-Agent Systems", the EPSRC responsive mode research project "Methods for reliability and control for autonomous underwater vehicles", the EPSRC responsive mode research project "Verification of security protocols: a multiagent systems approach", and the BT-funded research project "verification of networked systems" (this project hosted by University College London). Until October 2009 he was PI at Imperial for the EU-IST Strep project Contract (Contract based e-Business System Engineering for robust, verifiable Cross-Organisational Business Applications) which was ranked as “excellent” in its final review in September of the same year. Before 2009 he served as PI for a number of other projects including 2 further EPSRC research projects.

2.2.5 Technische Universiteit Eindhoven (TU/e)

The Architecture of Information Systems (AIS) research group of the Technische Universiteit Eindhoven (TU/e) is participating in this project. TU/e is one of Europe's leading research universities in Engineering Science & Technology and contributes to the advancement of engineering science and the development of societal and technological innovations. The Architecture of Information Systems (AIS) research group at TU/e investigates methods, techniques and tools for the design and analysis of Process-Aware Information Systems (PAIS), i.e., systems that support business processes (workflows) inside and between organizations. Hence, the focus is not limited to information systems and their architecture, but also the modeling and analysis of business processes and the organizations they support. AIS’s mission is to be one of the worldwide leading research groups in process modeling and analysis, process mining, and PAIS technology. The group driven by the motto "Process Technology that Works" and is well-known for innovations that are highly original and applicable in real-life situations. AIS researchers are working on a wide range of topics including workflow management, process mining, simulation, Petri nets, business process management, process modeling, and process analysis. This resulted not only in landmark publications but also in software products and true impact in industry.

The group is known for its work on process analysis (in particular process mining and verification) and workflow management. For example, the notion of soundness and corresponding workflow verification tools has been widely adopted by both academics and practitioners. The Workflow Patterns initiative (collaboration with QUT) has influenced several standardization processes and has become a standard tool for the selection of WFM/BPM technology. In the last decade hundreds of patterns have been collected and these are distributed via the website www.workflowpatterns.com (the most visited website in the workflow area for many years). More recently, the work of AIS on process mining has had a significant impact on the BPM field. The ProM tool, developed by members of the AIS group, has been applied in dozens of companies. Many of the ideas in ProM have been re-implemented in commercial tools such as Protos, BPM - one, Futura Reflect, ARIS PPM, etc. In fact, the work on process mining done at TU/e is seen as one of the most important innovations in the BPM field by Gartner.

Role attributed in ACSI

TU/e's main contribution will be to the WP on monitoring and mining (WP3). Here the experience and know-how are obvious. TU/e will also be involved in artifact-centric modeling. Here the know-how from Proclets will be used to revolve modeling problems. The link between both is also interesting. For example, the discovery of artifact-centric processes is a challenging and highly relevant problem.

Key personnel

Prof.dr.ir. Wil van der Aalst is a full professor of Information Systems at the Technische Universiteit Eindhoven (TU/e) having a position in both the Department of Mathematics and Computer Science and the Department of Technology Management. Currently he is also an adjunct professor at Queensland University of Technology (QUT) working within the BPM group there. His research interests include workflow management, process mining, Petri nets, business process management, process modeling, and process analysis. Wil van der Aalst has published more than 115 journal papers, 15 books (as author or editor), 230 refereed conference/workshop publications, and 40 book chapters. Many of his papers are highly cited (he has an H-index of more than 64 according to Google Scholar, making him the Dutch computer scientist with the highest H-index) and his ideas have influenced researchers, software developers, and standardization committees working on process support. He has been a co-chair of many leading conferences (BPM, ATPN, SCC, etc.) and is also editor/member of the editorial board of several journals, including the Distributed and Parallel Databases, the International Journal of Business Process Integration and Management, the International Journal on Enterprise Modeling and Information Systems Architectures, Computers in Industry, Business & Information Systems Engineering, IEEE Transactions on Services Computing, Lecture Notes in Business Information Processing, and Transactions on Petri Nets and Other Models of Concurrency. He is also a member of the Royal Holland Society of Sciences and Humanities (Koninklijke Hollandsche Maatschappij der Wetenschappen).

Dr. Boudewijn van Dongen is an assistant professor in the Information Systems group of the Department of Mathematics and Computer Science of Eindhoven University of Technology, Eindhoven, The Netherlands. He received his Ph.D. in 2007, after successfully defending his thesis entitled "Process Mining and Verification". Currently, his research interests extend from process mining and process verification to supporting flexible processes and visualization of research results. Furthermore, he plays an important role in the development of the open-source process mining framework ProM, freely available from www.processmining.org.

2.2.6 Tartu Ulikool (UT)

With around 18000 students and 1400 academic staff, University of Tartu (UT) is the largest university in Estonia. Founded in 1632, UT is member of the Coimbra Group, uniting reputable European research universities of long-standing traditions.

UT's Institute of Computer Science (CS) has over 400 students (incl. 35 doctoral students) and 34 academic staff. The institute conducts research in programming languages, software engineering, distributed systems, cryptography, bioinformatics, data mining, and language technology.

The Institute is a key research partner in the Estonian Center of Excellence in Computer Science and in the Software Technology and Applications Competence Centre – an R&D centre that involves 10 industry partners and conducts industry-driven research projects in the fields of services engineering and data mining. The institute is involved in four FP6/FP7 projects as well as one EUREKA project in the field of Software-as-a-Service and cloud computing.

Role attributed in ACSI

UT will lead WP3 (observation-based techniques and tools). In this WP, UT will leverage its expertise in two areas: (i) runtime adaptation of Web services; and (ii) conformance checking of service behavior (work done in collaboration with TU/e). UT will also contribute in WP1 to the definition of the artifact-centric process modeling framework. UT has previously participated in the design and implementation of the FlexConnect object-centric process design toolset and will carry forward this expertise into the project.

Key personnel

Prof. Dr. Marlon Dumas is holder of the Swedbank Professorship of Software Engineering at University of Tartu since 2007. Prior to this appointment, he was Associate Professor at Queensland University of Technology (QUT), Australia. During his time at QUT, he was Chief Investigator in 6 research projects funded by the Australian Research Council. He also received a prestigious Queensland Government Fellowship between 2004 and 2007 to undertake research on service-oriented software architectures in collaboration with SAP AG. Prof. Dumas has co-authored over 120 publications in international conferences and journals, 2 U.S. patents (granted), 3 other U.S. patent applications, and a textbook on Process-Aware Information Systems. His research papers have been cited over 4500 times according to Google scholar and have earned 4 best paper or best student paper awards at international conferences (CEC’2009, FASE’2007, EEE’2006, EUC’2005). He served as program chair of the 2007 International Conference on Business Process Management and demonstrations chair of the 2008 International Conference on Service-Oriented Computing. He received his doctoral degree from University of Grenoble in 2000.

Dr. Peep Küngas is currently researcher (since 2009) at University of Tartu where he carries out research on services ecosystems and modern Internet technologies. From 2007 to 2009 he was a post-doctoral fellow at the Department of Microelectronics and Information Technology of the Royal Institute of Technology, Kista, Sweden where he was involved in an EU project RoboSWARM and provided his expertise in knowledge systems and Web services. Peep Küngas received PhD in Information and Intelligence Science from Norwegian University of Science and Technology, Trondheim, Norway in August 2006. MSc and BSc in Computer and Systems Engineering were acquired from Tallinn University of Technology, Tallinn, Estonia respectively in 2002 and 2000. Peep has expert knowledge and experience in the Semantic

Web, semantic interoperability, formal methods, automated theorem proving, knowledge engineering, automated Web services annotation, intelligent distributed systems, analysis of service networks and large-scale federated information systems. Peep has served as a Program Committee member and reviewer of a number of scientific forums. Additionally he has published about 50 scientific peer-reviewed and more than 20 popular scientific papers. In parallel to his academic career, Peep maintains an online marketplace for XML Web services (SOATrader.com). This marketplace publishes, analyses, and brokers access to over 10000 publicly accessible Web services and aims at bringing together a community of Web services providers, consumers and developers for tight collaboration.

2.2.7 Indra Software Labs SLU (Indra)

INDRA Software Labs is part of the Indra Group, established in July 2004. With the integration of companies Azertia and Soluciona in 2006, Indra has grown to become the leading company in solutions and services involving high technology content in Spain and one of the main European and Latin-American companies. Indra is the premier information technology company in Spain and a leading IT multinational in Europe and Latin America. It is ranked as the second European company in its sector according to stock market capitalisation, and it is one of the top three Spanish companies with significant investment in R&D. In 2007, revenues exceeded € 2.167 M, of which a third came from the international market. Indra has become an outstanding reference on the markets on which it operates, both nationally and internationally. With references in more than 90 countries on the five continents, approximately a third of the company’s annual revenues come from the international markets. Indra is organised around six vertical markets: Defence and Security; Transport and Traffic; Energy and Industry; Telecom and Media; Finance and Insurance and Public Administration and Healthcare. The company employs more than 28.000 professionals.

Indra Software Labs consists of several technology development centres, whose objective is the development of new products and innovative services specialized in different areas. Products developed by Indra Software Labs have their own commercial operability, and have the support and the experience of Indra Group. Indra Software Labs works like a unique virtual centre providing services of software development in more than 25 countries, 24 hours a day, 7 days a week, which is one of its best actives. A local “offshore/near shore” development model, the continuous bet for R&D and the compromise with the generation of high quality technology products conform to the current vision of Indra SL. The technology development centres of Indra Software Labs located in Madrid, Badajoz, Malaga, Salamanca, Ciudad Real, Lérida, La Coruña, Bratislava, Ciudad de Panamá, Buenos Aires and Manila, are capable of producing global personalized services, developing customized and outsourcing solutions, as well as internal software developments. More than 1,500 professionals develop innovative products and specialized services for a variety of sectors. The collaboration with our technological partners and the internal quality requirements, given by the adoption of several ISO certificates, provides a guaranteed standard of excellence for projects carried out by Indra’s development centres.

<u>Role attributed in ACSI</u>

As ACSI industrial partner, Indra will develop use cases close to real requirements where the artifact centric approach and interoperation hub may solve actual problems and ease actual developments as well as studying the feasibility of the results of the project. Indra will lead WP5 (Validation) and will have a major role in the ACSI Hub System development (WP4).

Key personnel

Daniel Martínez Maqueda studied Software Engineer at Universidad Complutense de Madrid and specialised in AI techniques like automated learning and genetic algorithms. Daniel has broad experience in R&D field working for Telefonica I+D and nowadays working in Indra Software Labs as R&D Engineer. In Indra his main role is the execution of R&D national and international projects working in both managing and development tasks specially in SaaS and new web paradigms.

David Toribio Gómez studied Sofware Engineer at Universidad Autónoma de Madrid directing his studies in security technologies. He worked in Telefónica I+D in several national R&D projects related to advance biometry and information security techniques. In Indra he's been responsible of several national R&D projects related to home automation and DNIe developments from the security point of view.

2.2.8 Collibra NV/SA (Collibra)

Collibra was founded in June 2008 as a spin-off from the Semantics Technology and Applications Research Laboratory (STARLab) at the Vrije Universiteit Brussel (VUB). VUB STARLab is a centre of excellence in semantic technology research. This lab validated its methods and tools in many European (FP6 and 7) and national research projects.

Collibra was founded on a vision that business ecosystems have evolved from rigid value chains towards agile, globalized value networks. In this reality, the need for the exchange, integration and understanding of heterogeneous information has become of strategic and competitive importance. Despite its technological success story, the ICT enterprise is hampered by a rampant growth of metadata and a massive dump of scattered legacy data.

Collibra is convinced that architecting the next-generation enterprise requires a paradigm shift that goes beyond technical excel. Main drivers behind the dynamics of this so-called Enterprise 2.0 are the strategic and meaningful reconciliation of disparate data assets by describing them with so-called business semantics, and the pervasiveness of these processes in the daily workflows of business and IT. In other words, Collibra envisions semantics as the key enabling technology for governing and delivering heterogeneous information assets across the extended enterprise. Collibra's mission is to provide our customers with solutions for semantic data integration and metadata management. Currently Collibra has several customers in government and industry.

Role attributed in ACSI

Collibra's main contribution will be the development of the business semantics layer, and further extension of its techniques for semantic data integration (involvement in WP1 and WP2). Collibra will be deeply involved in WP5, validation of the ACSI Hub System, and also in the hub system development (WP4).

Key personnel

Dr. Pieter De Leenheer is assistant professor at VU University Amsterdam in the Business, Web and Media group (since Nov 09). He is co-founder of Collibra SA in Brussels. From 2002-2009, Pieter was senior scientist at VUB STARLab, and lecturer at the same university. Pieter holds a PhD in computer science on community-based ontology evolution, and a MSc in principle computer science. From 2004-2006, Pieter represented VUB STARLab in the European IP project DIP dealing with the integration of data, information, and processes. He worked mainly on the ontology versioning method and tool developed for the WSML framework. DIP ended with excellence. From 2005-2007, he was involved in the European Leonardo Da Vinci project were ontologies are built for meaningful competency-centric human resource management.

Pieter authored more than 30 publications in various books, international journals and conferences, among which he co-edited the Springer book "Ontology Management for the Semantic Web". He gives master lectures including Database Theory, (Web) Information Systems, and Semantic Web languages. He is member of ACM and IEEE.

Stijn Christiaens is co-founder and COO of Collibra. He has been an R&D engineer in the supply-chain and warehouse management industry and a researcher at the Semantics Technology and Applications Research Laboratory (STARLab) at the Vrije Universiteit Brussel in Brussels, Belgium. He contributed to several national and European projects on ontology engineering. Stijn holds an MSc in Industrial Engineering (in IT) from the Katholieke Hogeschool Gent (KiHo), an MSc in Artificial Intelligence from the Katholieke Universiteit Leuven, a degree in Project Management (PMI) from the European University College and a Postgraduate in Industrial Corporate Governance at the European University College.

B 2.3 Consortium as a whole

2.3.1 Expertise of partners

The project team includes world-class researchers in all of the key technical areas needed for this research, including experts on artifact-centric business processes, verification, ontologies, process mining, services architectures, and business process management, as described in Table 2-3 below. The consortium team comprises researchers who are not only outstanding scientists in their domains, but also prolific authors who have published dozens of first-class scientific articles and books, and presented their work at various leading conferences around the globe. To see a partial list of these works, please see the Partner References section at the end of this document.

Table 2-3: Partner expertise

Partner	Background/expertise
IBM	IBM brings to ACSI its extensive skills and experience in tool integration frameworks, the artifact-centric approach, Service Architectures, and Business Process Management. The IBM team for the ACSI proposal are members of the larger IBM Research team that pioneered the artifact-centric approach to Business Process Management in the early 2000's. Since 2005 that team has been applying artifact-centric methods and tools with clients both inside and external to IBM. The IBM ACSI team recently introduced the preliminary version of artifact-centric interoperation hub framework, with an interest in extending the application of artifacts to the realm of service interoperation. The IBM ACSI team has made extensive research contributions in database theory, database programming languages, business process and workflow management, and more recently web services. The IBM ACSI team has deep experience in business process

Partner	Background/expertise
	management, with application to both medium- and large-scale industry sectors, including telecommunications, finance and large-scale supply chains.
Uniroma1	The research group Uniroma1, is one of the best known research groups in Description Logic (DL) research. Its contributions in this field have been numerous and crucial, including, DLs for conceptual modeling in databases, formal and computational analysis of expressive DLs (which became the basis for the W3C Ontology language OWL), studies on conjunctive query answering in DLs, and the introduction of ontology-based data access and integration. The group has also contributed very deeply to the research on reasoning about actions in AI. Prof. De Giacomo, was one of the inventors of the ConGolog high level language based on situation calculus, developed together with Hector Levesque and Yves Lesperance, and then of IndiGolog which focuses on mixing deliberation and execution. The group has worked at Service/Process Composition and Integration, developing one of the most prominent theoretical models for composition of conversational services: the "Roman Model". The group has also contributed very profoundly to Database research. Indeed the group is considered one of the best research groups worldwide in data integration, and management of semistructured data (the latter investigated with Moshe Vardi). In particular, the group contributed to the formal definition of the data integration problem. A famous tutorial by Maurizio Lenzerini at PODS on data integration is one of the most cited references on this subject.
UT	1. Process modeling. Members of the group have previous experience in the definition of a business object-centric process modeling language. This expertise will be used in the definition of the artifact-centric process modeling framework. 2. Service composition and adapter synthesis: Members of the group have previous expertise in models for service composition both using imperative approaches (statecharts) and logic-based approaches (linear logic). Also, the group has expertise in synthesis of Web service adapters from mapping rules. This expertise will be applied in the WP on Verification and Synthesis. 3. Service network analysis. The team has expertise in building service network models from repositories of semantically-annotated service descriptions, and using these network models to predict ownership of data (which data is managed by which service?) and data redundancy. This expertise will further developed and exploited in the WP on Monitoring and Mining in order to detect possible groupings of services into artifacts and to determine which data attributes of an artifact are read or modified by a given service.
Imperial	Multi-agent Systems. The group has long expertise with formalisms and implementations of multi-agent systems, including negotiation frameworks and implementations, auction systems, and web services. Logics for Distributed Systems and Multi-Agent Systems. The team has

Partner	Background/expertise
	long standing expertise in the development and study of modal logics for the specification of multi-agent systems. This includes contributions in temporal logic, epistemic logic, deontic logic, and alternating time logic. Verification of multi-agent systems. The group has pioneering expertise in model checking multi-agent systems. This includes symbolic approaches based on ordered-binary-decision diagrams as well as SAT-based, including bounded model checking and unbounded model checking. The group authored and maintains MCMAS, the leading bdd-based symbolic model checker for multi-agent systems. Recently this technology has been employed for the verification of web services as part of EC project CONTRACT. Further experience. The group is active on a number of contiguous areas to the core of the project that may become of interest to the consortium. This includes work on symbolic monitoring of web services, security protocols, and diagnosability via model checking.
FUB	FUB brings to the project expertise in description logics and ontology languages optimized for ontology-based data access and integration. It has contributed to the development of the QuOnto system, and has developed novel algorithms for the efficient processing of mapping assertions when answering queries over ontologies. It has expertise in reasoning in the presence of ontologies combined with rules. Moverover, it brings expertise in temporal conceptual modeling and in the combination of description logics and temporal logics.
TU/e	The TU/e will contribute its expertise in the following areas 1. Process mining: The TU/e is the leading research group in process mining. There is a lot of experience in developing process mining algorithms and their practical application. Lots of this knowledge is embedded in the ProM tool. This open-source process mining tool was initially developed at TU/e and currently many other groups are joining this effort. The main contribution of TU/e to this project is in the field of process mining. 2. Service/process modeling: The group is one of the largest Petri net groups in the world. The groups has been working on the modeling services and workflows in terms of Petri nets (workflow nets, open nets, etc.). The group also worked on Proclets. These can be seen as a precursor of today's artifact-centric models. 3. Verification: TU/e has developed Petri-net-based verification techniques. Using tools such as Woflan and ProM complex verification questions can be tackled.
Collibra	Collibra is commercialising more than 100 person years in application-oriented research in ontology management tools and techniques (at VUB STARLab (see previous experience) that are grounded in natural information analysis methods (e.g., NIAM, ORM). The main expertise and technology that Collibra brings forward as co-foundation of this project is

Partner	Background/expertise
	twofold: Business Semantics Management, and Semantic Data Integration. Collibra's core products provide strong solutions that validate ontology engineering prototypes and methods for a broad range of varying domains, as demonstrated by the following examples. In the human resources domain, at CoDrive and Prolix, Collibra provided ontology evolution support for business processes to enable organisations to improve the competencies of their employees, thereby responding faster to continuous changes in business requirements. In the Web service domain, at DIP, Collibra contributed to the development and validation ontology versioning support for Semantic Web Services as a scaleable and cost-effective solution to the Web data and process integration problem. In the financial forensics domain, at FFPoirot, they developed an ontology engineering methodology. This resource (= ontology), or parts of it, could be commercially exploited as a set of Semantic Web services. The FFPoirot project was coordinated by STARLab and finalised with excellence (cf. EC report on the website) and hence presented at the Communicating European Research fair. Currently, Collibra has two patents pending for its technology.
INDRA	Information and Communication systems. Software as a Service, Cloud computing and web services. The Indra Software Labs' main client is Union Fenosa, which is one of the largest energy companies in Spain. In fact, Indra Software Labs is the name of the former Soluziona company, the former consulting and development company owned by Union Fenosa for software development. Indra works with the most important water, gas, and electricity utilities in the country (including Red Eléctrica, OMEL, ENAGAS, and Unión Fenosa). This experience provides Indra with an in-depth and updated understanding of the standards and trends of these sectors, which is brought into play in our consulting and integration projects for the market agents. Indra's global experience and profound understanding of the circumstances and needs of the companies involved in these sectors put the company in a position to design systems and solutions that, through an optimum combination of technologies, processes, and tools, along with the strategies of each entity, ensure the success of all their clients in this new economic environment, allowing them to reduce costs and continuously improve their competitiveness and the service provided to the users.

2.3.2 Balance and Coverage

The partners of the ACSI consortium are key European players with recognized leadership, proven experience, and complementary skills in research and development in the areas relevant to the ACSI proposal. The consortium presented a balanced combination of small and large companies and scientific institutes from various European countries, with complementary capabilities. Furthermore, two of the industrial partners (one of them an SME), in addition to their research contributions, will be responsible for evaluating ACSI Hub

System developed via case studies in two different domains. The overlaps and complementarities among expertises of the ACSI participants ensure that, in every workpacakge, there will be a fruitful cooperation among the partners.

The different activities of research, development, and exploitation will be shared among the partners according to their relative strengths and interests.

2.3.3 Partners by country of origin

The project aims to approach as wide target audience as possible. The geographic diversity of the consortium and various languages the partners speak as well as various cultural origins will strengthen the capabilities of the consortium. The multilingual nature of the consortium overcoming the language barrier enables wider approach to the national data (online/offline sources, national projects and initiatives) and expands the vision of the project.

Table 2-4: Partners by country of origin

Participant org, short name	Country
IBM	Israel
Uniroma1	Italy
FUB	Italy
Imperial	UK
TU/e	Netherlands
UT	Esthonia
Indra	Spain
Collibra	Belgium

2.3.4 Experience in EC projects

Table 2-5: Experience in EC projects

Partner	Project(s)
IBM	IBM Israel has been active in framework programs since FP4. With fifteen running projects, three of them in the coordination role, IBM is an Israeli leader in FP participation. Three projects currently coordinated by IBM are: RESERVOIR (FP7 IP project), +Spaces (FP7 STReP), and ModelPlex (FP6 IP project). This is in addition to PROSYD, SAPIR, and SHADOWS, FP6 STREP's successfully concluded. Among the projects with significant IBM Israel contribution: CASPAR, EuResist and SPEEDS (FP6) HyperGenes Impact and HERMES projects (FP7).

Partner	Project(s)
Uniroma1	The group has been involved in many national and European projects, including the EU funded projects Esprit 22469 DWQ, IST-2001-34825 SEWASIE (www.sewasie.org), IST-2001-33570 INFOMIX, IST-508011 INTEROP, FP6-7603 TONES, FP6-5-034749 WORKPAD, and FP7-224332 SM4All, and IFP7-ICT-2007-C-FET-Open VISMASTER CA.
FUB	Members of FUB have been involved in several EU projects, starting from FP5, as principal investigators, co-investigators, and coordinators. Among these, we mention: TONES (FP6-7603) as coordinator, KnowledgeWeb (IST-2004-507482) was an FP6 Network of Excellence, and OntoRule (FP7-231875).
Imperial	EU Marie Curie FP7 project "First-order Modal Logics for the Specification and Verification of Multi-Agent Systems and STREP project Contract (Contract based e-Business System Engineering for robust, verifiable Cross-Organisational Business Applications) which was ranked as “excellent” in its final review.
TU/E	Semantics Utilised for Process management within and between EnteRprises (SUPER) [FP6-IP 2007-2009], Cross-Organisational Workflow Formation and Enactment (Crosswork) IST-507590, Interoperability research for networked enterprises applications and software (Interop) IST-508011
UT	The TARTU institute coordinates an EU FP6 project (COBRED – bioinformatics) and it is partner in 3 other FP6/FP7 projects as well as one EUREKA project in the field of Software-as-a-Service and cloud computing SITIO (Semantic Business Processes Based on Software-as-a-Service and Cloud Computing).
Indra	Indra Software Labs worked in previous EU projects related to accessibility and assistive learning: ENABLED - Enhanced Network Accessibility for the Blind and Visually Impaired and EU4ALL - European Unified Aproach for Assistive Lifelong Learning.
Collibra	Collibra relies on the broad experience of participating and coordinating EU projects it inherits as spin-off company from Vrije Universiteit Brussel, and VUB STARLab in particular. Relevant projects in which STARLab participated include: Knowledge Web, CoDrive, FF Poirot, Prolix, and TAS3.

2.3.5 Sub-contracting

No relevant in the ACSI proposal

2.3.6 Other countries

No relevant in the ACSI proposal

2.3.7 Additional partners

No relevant in the ACSI proposal

B3. IMPACT

B 3.1 Strategic impact

3.1.1 Expected impact – compliance table

The ACSI project addresses strategic *Objective1.2: Internet of Services, Software and Virtualisation b) Highly Innovative Service / Software Engineering*. ACSI expected research and technological outcomes closely match the expected outcomes of this objective, and therefore should provide the results sought. In what follows, we list the expected impacts as they appear in the call, and explain how ACSI will contribute to deliver these impacts.

Table 3-1 Expected Impact – compliance table

Expected impact	ACSI contribution
"A major contribution to the Future Internet in terms of service development, management and interoperability"	Out-sourcing, globalization and automation of business processes continue to increase. However, today, there is no effective, flexible, scalable, widely used, and principled approach to enable *service collaborations*, that is, the interoperation of services across enterprise boundaries in support of shared (business) goals. This is a major roadblock preventing substantial automation of these kinds of collaboration, and more broadly, the design, deployment, and operation of innovative value nets. **The ACSI project is aimed directly at filling this vacuum.** The ACSI interoperation hub framework provides a new, general-purpose, structured, and open approach to enable service collaborations in essentially all industry sectors. The underlying notion of interoperation hub is drawn from successful examples such as the EasyChair conference management system, and SaaS offerings such as SalesForce.com The ACSI project will extend that core notion in several important ways. Perhaps most importantly, the ACSI hub framework is based on the use of "dynamic artifacts" (also known as "business artifacts") to model the business processes underlying a hub, making it easy and intuitive for designers to create new hubs, and service designers to understand the service collaborations they are joining. **Dynamic artifacts have a proven track record** for enabling a top-down, end-to-end view of business operations in a single-enterprise setting and enabling substantially more efficient design and deployment of business processes (see Subsection 1.1.1 for details on dynamic artifacts), and the ACSI research will bring those benefits to the realm of service interoperation. *Importantly, the services participating with an interoperation hub do not have to be artifact-centric; they can be arbitrary SOA services (including legacy services with a SOA wrapper).* The ACSI interoperation hub framework has a formal basis, and

Expected impact	ACSI contribution
	includes sophisticated tools to assist with conformance testing, verification, automatic code generation, and schema evolution. **The ACSI interoperation hub framework will provide a paradigm shift in the way that services, and more generally enterprises, can work together.** This shift will enable dramatic efficiency gains and foster substantial innovation in the conceptualization of future service collaborations and value nets.
"Deep technological advances in software/service engineering. New software technologies for improving scalability and predictability of distributed systems,"	The use of interoperation hubs, as developed in the ACSI project, will advance the science and technology in software and service engineering along four key dimensions. (1) First is the substantial development of the new framework for service interoperation, based on interoperation hubs, as an alternative to the now popular orchestration and choreography. **The ACSI interoperation hub framework has been designed specifically for an open network with long-tail characteristics, whereas orchestration and choreography are better suited for use-once kinds of deployments.** In particular, ACSI interoperation hubs enable scalability in two fundamental ways. First, an interoperation hub creates what is essentially a *de facto* standard for a category of service collaborations, and the individual services will have to adapt to it. This spreads the labor, enabling a hub to support a large number of disjoint service collaborations with similar goals. Second, the logic underlying the hubs is specified using dynamic artifact types, which have been proven, at least in the single-enterprise setting, to gracefully support rich variation. Thus, the same hub will be able to support multiple disjoint service collaborations of a given kind, even if they all have different policies and characteristics. The ACSI project will establish the framework, develop theoretical and pragmatic foundations, and build the ACSI Hub System for designing, deploying, and running large-scale service collaborations based on this framework. (2) ACSI interoperation hubs will take advantage of dynamic artifacts (see Subsection 1.1.1). Because they focus on the end-to-end lifecycle of key conceptual entities that pass through a business processes, **dynamic artifacts can aid substantially in the management of "hand-offs" of data and processing between services and organizations.** The ACSI research will develop a formal foundation for understanding this construct, which

Expected impact	ACSI contribution
	up until now has been studied and applied largely from the method and systems perspective. **Dynamic artifacts are based on a novel blend of data and process, enabling the rigorous study and application of these two central notions within a single conceptual building block.** This holds the promise of revealing new principles for understanding, managing, and implementing services. The formal foundation will introduce the notion of Knowledge and Action Base (KAB), which is analogous to the notion of ontology but is focused on rigorous modeling of both data and activity. This will be applied to develop a formal foundation for managing the evolution of artifact schemas (and interoperation hubs), including managing in-flight enactments during schema changes. The KAB reasoning tools developed as part of the ACSI research will be a significant technology advance over the current state of the art. (3) **The ACSI research will substantially extend current verification and synthesis techniques, to incorporate data along with process.** This is a notoriously difficult challenge, because the incorporation of data makes the underlying state space infinite, and so sophisticated techniques will need to be devised, to find abstractions that are rich enough to have general usefulness but restricted enough so that verification is feasible. The particular way that dynamic artifacts blend data and process offers promising starting points for this research. The verification and synthesis tools developed in this part of the ACSI research will open an entirely new area in which verification can be usefully applied. (4) **The ACSI research will develop the next generation of process mining research, by generalizing it to work with data along with process.** Using this, the ACSI research will develop capabilities for conformance checking between conventional services and artifact-based interoperation hubs, for auto-generation of adapters to simplify service on-boarding, and for mining (i.e., reverse-engineering) the possibly implicit dynamic artifacts of a hub or service. This process mining technology, together with the formal verification and synthesis technology, will substantially simplify the issues around service on-boarding and maintenance for service collaborations and value nets. By complementing the artifact-based runtime environment with cloud computing infrastructure, delivery of scalable services with high performance and availability in

Expected impact	ACSI contribution
	distributed settings becomes reality. The resulting techniques and tools will yield substantial improvements in the creation, maintenance, and predictability of service collaborations in distributed contexts.
"A more competitive environment including infrastructure operators moved up the value chain with innovative service offerings on scalable infrastructure."	The interoperation hub framework provides rich new opportunities for infrastructure operators, and other businesses that want to host these hubs. All of these businesses can create and/or host an interoperation hub for a given industry sector or niche. **This enables the kind of infrastructure being provided to be semantically richer than simple bit transport, and to contribute much more directly to the business-level aspects of enabling communication between individuals and enterprises.** As demonstrated by IBM in the context of single-enterprise solutions, the use of the artifact-centric approach fundamentally reduces the barriers between business analysts and IT in design and creation of business processes. We expect the same reduction in design and deployment of hubs, enabling more creativity in the conception and design of new kinds of service collaborations while lowering the costs. Indeed, the creation of interoperation hubs will offer new opportunities for infrastructure providers and others to facilitate the creation of new, innovative value chains. The ACSI interoperation hub framework, when fully mature, will support efficient interoperation in industry segments such as government, healthcare, supply chain, e-commerce, and large-scale manufacture (e.g., Airbus).
"Lowered barriers for service providers, in particular SMEs, to develop services through standardised open (source) platforms and interfaces"	Interoperation hubs will serve as *de facto* standards that can provide a structure within which collections of services (and enterprises) can collaborate to form service collaborations, and more generally, value chains. The use of artifacts in the hubs as the underlying model for business processes helps to ensure that the hub's logic can accommodate wide variations, and is easy to comprehend for the enterprises that want to participate in service collaborations. (Again, although the hubs themselves are based on dynamic artifacts, the participating services can be arbitrary SOA services.) **To simplify and reduce the costs of joining a service collaboration, the ACSI project will develop several targeted techniques and tools**, including auto-generation of adapters, automatic conformance testing, and static and dynamic verification capabilities. These techniques and tools will be geared towards open and long-tail contexts. The ACSI research will develop a declarative approach for specifying the

Expected impact	ACSI contribution
	dynamic artifacts, with the goal of making them as simple and lightweight to use and connect to as today's spreadsheet software is for data representation and computation. Finally, interoperation hubs can be placed into a cloud, and can provide data storage and task executions on behalf of the participating services, so the IT costs to SMEs can be based on usage.
"Massive uptake of high-added value services through innovative service front ends and a higher user empowerment and more advanced and dynamic online communities through platforms enabling third party generated services".	While the focus of ACSI is not on service front ends, the ACSI interoperation hub framework will make it faster and much less expensive to create new value networks that can bring together collections of services in support of innovative user-facing services (and service front ends). Interoperation hubs make it simple to create the artifact schema for new categories of service collaborations, and enable third party generated services to quickly join, and richly collaborate with, service collaborations.
"A strengthened industry in Europe for software, software services and Web services, offering a greater number of more reliable and affordable services, enabled by flexible and resilient platforms for software/service engineering, design, development, management and interoperability. Technologies tailored to meet key societal and economical needs"	The ACSI interoperation hub framework provides a flexible and resilient platform for service interoperation. **An ACSI interoperation hub can provide the anchor for a category of service collaborations, and provide a clear structure within which participating services can fit.** This overcomes a significant challenge that has been facing the software/Web services industry up to this point, because it is hard to create a targeted web service in the absence of an existing value net. If a small-scale value net is launched around an ACSI interoperation hub, then it can grow and expand in functional capabilities, because of the ACSI support for customization and evolution of the artifact schemas hub. The approach includes techniques and tools to enable rich variability and support the "long-tail". This should help support the creation of new businesses and enable SMEs to launch innovative niche services. The pilot projects will demonstrate the utility of the approach in additional areas relating to key societal needs, namely government services and energy creation (including small-scale) and distribution. Importantly, **much of the software prototyping for this project will be performed in an open-source manner**, making it easier for European software companies to quickly leverage the ACSI technology.

3.1.2 *Technological impact*

The ACSI project will perform pre-competitive research into the science and technology underlying a new and very promising approach for enabling rapid, flexible, innovative service interoperation. **The ACSI interoperation hub framework offers an approach to populating the web with semantically rich building blocks, that services can cluster around to create a broad variety of service collaborations and value networks.** ACSI interoperation hubs can be thought of as a qualitative advance over successful "hubs" such as the EasyChair conference management system and Salesforce.com, along two dimensions. First, because they are based on dynamic artifacts, which combine process and data into an holistic building block, it will be dramatically easier than with conventional techniques to develop ACSI interoperation hubs for a broad array of application domains that can benefit from collaborations among multiple services. Second, because ACSI will provide tools to simplify on-boarding into an ACSI-enabled service collaboration, the idea of value nets consisting of medium and large numbers of automated and/or manual-assisted services can become a reality.

Through scientific advances, substantial prototype development, and pilot testing, the ACSI project will firmly establish the structures, techniques, and base technologies upon which industrial strength interoperation hub environments can be developed. Most of the ACSI research results will be combined into the ACSI Hub System, which will include a comprehensive Design Workbench for designing, deploying, maintaining, and evolving interoperation hubs. This will be complemented with the Hub Engine, that can support numerous distinct service collaborations guided by the hub's artifact schema, including capabilities for coordinating each service collaboration, on-boarding new collaborations and new services participating in them, and providing in-hub data storage and task execution for those services that desire it. This infrastructure will include ***access controls***, to ensure that a participating service can "see" only an appropriate subset of all the artifact instances being managed by the hub, can "see" only an appropriate portion of each artifact instance (e.g., only some of the attributes and only some portions of the full lifecycle), and can execute only an appropriate subset of the tasks and lifecycle transitions for these artifact instances. The framework will also incorporate three substantial families of scientific advances with supporting prototype tools. First is a suite of capabilities stemming from the ***adaptation of process mining*** techniques to the artifact-centric context, enabling conformance checking, on-the-fly adapter generation, and mining (i.e., reverse-engineering) of dynamic artifact types for hubs and participating services. Second is a suite of ***design-time and run-time verification tools***, to aid in both the design and on-boarding processes, ***and also automatic synthesis tools*** to further aid in the design of the artifact schemas in interoperation hubs. All of these scientific advances will represent substantial advances over the current state of the art, because they will incorporate data along with process as a first-class citizen. In contrast, the third scientific advance will create a substantial generalization of a data-focused technology, namely ontologies and ontological reasoners. This advance will develop a theory of ***Knowledge and Action Bases* (*KABs*)** that can form a semantic layer above dynamic artifacts. KABs will enable reasoning about, and comparing, artifact schemas, in a manner analogous to how ontologies enable reasoning about, and comparing, database schemas. This will be used in the research to develop tooling in support of evolving the artifact schemas of interoperation hubs. Finally, the integration of all of these capabilities in the ACSI Hub System will be repeatedly evaluated against two representative use cases (one from Energy Distribution Management and the other from EU government publications) with lessons learned incorporated into the ACSI Hub System.

The ACSI interoperation hub framework, in conjunction with the underlying ACSI artifact paradigm, provides a rich structure around which many subsequent scientific and technology advances can be made. While the ACSI research will establish the basis for artifact-centric interoperation hubs, and key scientific and technology results in support of their easy use and operation, there will be a continued need for research and technology development in several directions. Issues such as user-centric design and runtime optimizations, while touched upon by the ACSI research, will deserve further attention once the basic ACSI structure is established. The three scientific thrusts of the ACSI research each address very fundamental questions, stemming from the convergence of data and process found in dynamic artifacts. While ACSI will deliver practically useful results for the context of interoperation hubs, it will also leave many questions open for exploration from both scientific and technology perspectives. Finally, while the ACSI research will establish the structure and basic algorithms and tools for access controls, this will be a very real concern for many services participating with an interoperation hub, and so further research will be warranted there as well. By centring the ACSI research program in Europe, the EU will be able to take advantage both of a rapid industrialization of the ACSI technology, and also have a head start on the further development of technologies in support of flexible, scalable service interoperation.

3.1.3 Socio-economic benefits of the project

Interoperation of services, be they within an organization or across organizations, remains one of the dominant challenges in the Information Age. While hand-crafted service interoperations are now common, and provide key support for national and international infrastructures in such sectors as banking, telecommunications, energy, and law enforcement, these are often inflexible, and require significant effort to add new capabilities. But for many service collaborations, the interoperation between services is often accomplished through partially automated/partially manual processes, leading to massive inefficiencies, in terms of both labor expended and time lost to an overall process enactment. As noted in the European Interoperability Framework for Pan-European e-Government Services [http://ec.europa.eu/idabc/en/document/3473], "there is a growing awareness that the interoperability of national public ICT infrastructures is a precondition for a more service-oriented and competitive public sector." While the manual processing for the hand-off between services does afford flexibility, it often introduces an *ad hoc* characteristic, making audits and analysis more challenging, and evolution of the overall process expensive and again *ad hoc*.

The ACSI interoperation hub technology aims directly at the core challenges around service interoperation, namely a current lack of systematic structure for service collaborations, the need for variability and flexibility when supporting them, and the need for a principled approach to support the hand-off of process and data between services. The project aims to achieve dramatic savings over conventional approaches to service interoperation across several areas (design and deployment; on-boarding; day-to-day operation; maintenance and evolvability), while still enabling rich flexibility for the different service collaborations using a given interoperation hub. **The technology will be applicable in key challenge areas of societal importance, including government, energy, healthcare, supply chain logistics** (especially in industries such as food or heavy manufacture with deep upstream supply chains), **and heavy manufacture** (e.g., Airbus).

The ACSI interoperation hub technology is positioned to take advantage of the current trend towards Software as a Service (SaaS). According to [Gartner User Survey Analysis: Saas, Enterprise Application Markets Worldwide, October 2008], "nearly 90% of organisations [worldwide] expect to maintain or grow their SaaS usage, with more than one third transitioning from on-premises to SaaS". And according to the Gartner group [Software-as-a-Service Tops Industry Trends, whitepaper, 2009], "one third of all new software will be delivered by on-demand by 2010." This indicates that the IT consumers are willing to rely on SaaS's in general, and it is thus likely that they would be willing to rely on SaaS's based on the ACSI interoperation hub technology. The Saugatuck group has the thesis "that SaaS is gaining significant traction in Europe – led by high adoption rates in the UK, Benelux and Scandinavia" [Readying for Strong SaaS Market Growth in Europe, Saugatuck technology, research alert report 2008]. However, they report that "the biggest US-based firms such as Salesforce (i.e., CRM/SFA) and other cross-industry players such as Taleo (i.e., Talent Management) are starting to gain momentum with broader-based horizontal offerings. Where this scenario might vary is around core human resource and finance-targeted apps, where country-specific rules and regulatory mandates provide a competitive advantage for European-based players." **The ACSI interoperation hub framework, with its systematic approach to supporting flexible, scalable collaboration environments might help the EU software vendors gain ground against the US competitors. Further, the mechanisms incorporated into the ACSI framework to support rich variation within a single hub could be especially advantageous in industries, such as Human Resources, where there are significant differences from country to country.**

B 3.2 Plan for the use and dissemination of foreground

The ACSI project will allocate the following efforts to guarantee dissemination and exploitation of its results:

- **Dissemination**: the partners will disseminate the results within the ICT community, the academic research community, and the general public.
- **Exploitation**: the partners will strive to exploit ACSI results in the industrial community via several routes as described in section 3.2.2.

ACSI has allocated a specific WP (WP6) to coordinate all project dissemination and exploitation activities. The details of the planned efforts are provided in the following sections.

3.2.1 Dissemination

The ACSI project plans to apply two major dissemination strategies. Firstly, the project will have a web site for posting project activity and achievements that will be accessible to the public. Web site construction and management are part of WP6. Secondly, the members of the project will publish technical and scientific results at scientific events. These events are open to the public and attract the most relevant audience, thus increasing public awareness. ACSI will also take advantage of Web 2.0 techniques to increase awareness, e.g., twittering of news feeds when tool updates are put online or publications are accepted at conferences; blogging intermediate results and announcements; discussion fora; social networking environments (LinkedIn, Twitter); etc. As part of these events, the project will organize an annual workshop for both scientific and industrial communities.

More specifically, the results achieved by the project will be disseminated by the following means:

- *Web Site*: The project will have a dedicated web site within the *.eu* domain, which will comprise elements such as the project description, objectives of the project, partner profiles and expertise, regular information on the project's progress and achievements, and downloadable versions of all public documents generated by the project. The web site will also be used as a publicly accessible repository for any promotional or disseminative material stemming from the project results. The project will register the domain name for two years beyond the project end date.
- *Electronic Newsletter*: Starting early in the project, we will send out an electronic newsletter publicizing the project and its outputs and invite subscribers, in an effort to build up a database of interested parties, including potential end users. The initial database of subscribers will be the partners 'existing contacts' database. The electronic newsletter will be managed by project participants for online subscription-based broadcasts.
- *Participation in technical events* organized either by the European Commission or by other software engineering projects, making the project's results available to the ICT-related sectors in Europe.
- *Academic courses*: Research and academic members of ACSI will disseminate project results in advanced academic courses at universities and in other scientific seminars and tutorials.
- *Press releases and publications* will be prepared and distributed in such formats as leaflets, brochures, articles, press releases, and other dissemination material.
- *ACSI workshop:* The project will consider organizing workshops on the ACSI research topics for the international research community and potential interested parties (including industrial representatives and investors).
- *Papers in the scientific literature*: Academic and research project partners will publish their work in related academic journals. In particular, work will be published in journals covering the different specialization of partners. Partners of the ACSI project will target the following journals for publishing and disseminating the project outcomes:

  IEEE Transactions on Services Computing, Elsevier journal Computer for Industry, IEEE Intelligent Systems, ACM Transactions on Database Systems (TODS), ACM Transactions on Computational Logic (TOCL), Journal of Computer and System Sciences (JCSS), Information and Computation (IC), Artificial Intelligence Journal (AIJ), Journal of Artificial Intelligence Research (JAIR), Information Systems (IS), Software and System Modeling, Journal of Cooperative Information Systems, ACM Transactions on Internet Technology.
- *Participation in conferences and workshops*: Partners of the ACSI project will target the following conferences for publishing and disseminating the project outcomes:

  Conference of Association for the Advancement of Artificial Intelligence (AAAI), International Conference on Knowledge Representation and Reasoning (KR), International Joint Conference on Artificial Intelligence (IJCAI), International Conference on Principles of Database Systems (PODS), International Conference on Database Theory (ICDT), International Conference on Very Large Data Bases (VLDB), International Conference on Service Oriented Computing (ICSOC), Business

Process Management Conference (BPM), International Conference on Autonomous Agents and Multiagent Systems (AAMAS), International Conference on Automated Planning and Scheduling (ICAPS), IEEE International Conference on Web Services (ICWS), IEEE International Conference on Services Computing (SCC), IFIP International Conference on Formal Techniques for Distributed Systems (FMOODS/FORTE), OnTheMove Federated Conferences & Workshops (OTM), Computer Aided Verification (CAV), International Semantic Web Conference (ISWC).

ACSI members are on the program and steering committees of some of these conferences and can thus stimulate events addressing service interoperation, such as panel discussions and keynotes. In the last third of the project, in one of these venues, we plan to organize a workshop in which we will give an overview of the project, discuss solutions and future challenges, and demonstrate the ACSI Hub System to the industrial and academic community (this is in addition to the annual ACSI workshop). These conferences, workshops, and journals are very focused on the research issues addressed in ACSI, therefore acceptance of papers implies originality of the work and enables fruitful discussions with a focused research community. Moreover, presenting ACSI at such venues guarantees a wide visibility of the project results.

3.2.2 Exploitation

The coordinator of the ACSI project will appoint an exploitation manager responsible for accelerating the introduction of ACSI results into the marketplace. We briefly outline below the initial exploitation directions and expectations for ACSI. Exploitation activities during the project lifetime consist of initial market analysis and elaboration strategy for further exploitation (as part of Deliverable 6.4 in the Dissemination and exploitation WP). Moreover, to promote exploitation, the ACSI project partners will present the ACSI Hub System at international conferences, as well as at dedicated events targeted towards the industrial community. Via these presentations, the results of the project will gain exposure to potential users, increasing the chances of commercial adoption of the technologies developed in ACSI.

ACSI will offer its partners and results adopters several routes for exploitation.

Direct Exploitation – In particular, the possibilities for further adoption of ACSI technology are realized through the potential integration in existing products through liaising partners of research institutes and directly by industry partners. The main industrial exploitation partners for the post-project phase are IBM, Collibra, and Indra.

IBM has a rich tradition of exploiting Research technology in its products. Indeed, several technologies developed by IBM's Haifa Research Lab play a major role in IBM's products. These products play a prominent role in the infrastructure of many organizations including hundreds of European companies. IBM will promote ACSI ideas by presenting ACSI technology to IBM executives and actively participating in development efforts to include ACSI technology in new products or offerings.

Indra is the premier information technology company in Spain and a leading IT multinational in Europe and Latin America. Indra has become an outstanding reference of the markets in which it operates, both nationally and internationally. With references in more than 90 countries on five continents, approximately a third of the company's annual revenues come from international markets. This solid presence in the international market with extensive and well-established business operations will prove critical to the successful exploitation of ACSI.

Indra aims to apply the results gained from the pilot in two energy client running projects (BDM and OMER).

Collibra is a regular sponsor of major international conferences both with an academic (OTM) and industrial (SemTech, Defrag, Data Management and Information Quality Conference, SOA Symposium Rotterdam) audience. The company intends to use these channels to promote its R&D activities and results, with an emphasis on how they fit in current commercial products. Collibra aims at incorporating ACSI results into their current product offering portfolio. Collibra aims at adopting the results from ACSI in their product suite. The Business Semantics Glossary can be extended with the results from the semantic layer development. The Business Semantics Studio can adopt the tools developed for verification. Finally, Collibra plans to incorporate the resulting methods and tools for the interoperation hub in their Business Semantics Enabler.

Special Interest Groups (SIGs) – SIGs will consist of researchers, companies, and authorities interested in the ACSI project. These might be interested in the implementation of the ACSI concept, application of its methodologies and, possibly, in early testing and use of its tools. SIGs will be populated using contacts of the key participating partners. SIG dissemination provides a unique vehicle for the project promotion. Activities will include participation in a joint meeting with related European projects, press releases, newsletters, brochures, demonstrations, and presentations to interested parties. These groups will consist of industrial parties interested in implementing the ACSI methodology and tools for testing, within the limits of the Special Interest Group. Activities will include the collection of the user requirements; live demonstrations of the developed software at project showrooms or online via the Virtual Project Showrooms (to obtain feedback for improvements); and presentations to the parties interested in commercial exploitation. For example, the University of Tartu will actively look for cooperation opportunities in the area of large-scale services ecosystems through its staff members at the Estonian National IT Interoperability Working Group. The working group is responsible for developing and enforcing usage of interoperability frameworks, aligned with IDABC activities, on a national scale in Estonia. Collibra is a member of the management committee of COST Action on agreement technologies (http://www.agreement-technologies.eu). Collibra is also affiliated with VU Amsterdam and maintains a strong connection with Tudor service science group in Luxemburg.

Education – Education also provides a modality for exploitation of ACSI results, through the strengthening of degree programs and other high-value courses to students. As an example, the University partners of Roma and Bolzano will develop a graduate course in the Faculty of Computer Science based on ACSI results. It is also expected that part of the results produced in the project will form part of the master's level curriculum in specialized courses at Imperial College. The University of Tartu aims at delivering lessons learned and novel techniques to the development of the national federated information system called X-Road, where the results of ACSI can be applied.

Industrial Training – Alongside educational institutes, ACSI will also be exploitable through the provision of professional and corporate courses to potential users, including public administrations. These training courses will be carried out by the consortium partners.

Special Advisory Board - In addition to the partners of the project, ACSI will strive to establish a Special Advisory Board (SAB). The members of this advisory board:

- Are NOT partners in the project and will NOT receive EC funding for research & development activities. Special Advisory Board members will also be asked to sign appropriate non-disclosure agreements before participating in any meetings.
- Are typically industrial organizations which will be interested in the project and will be willing to exchange views, provide the end-user perspective or a specific know-how, and generally support the project partners in their specific areas of expertise.

At this stage, we foresee that the roles of the SAB will be as follows:

- To contribute with their expertise to WP5 (Validation)
- To provide definitions of requirements and additional use cases to be validated
- To provide feedback, advice, and validation of the project developments
- To collaborate in the dissemination and exploitation activities of the project by advising how to maximize its potential impact.

Schedule for the elaboration of the exploitation plan - The following high level schedule for the elaboration of the exploitation plan is foreseen to facilitate the transition from the research phase to industrial exploitation:

- **6 months after project start -** Identification of market segments; set up long-term business strategy; address industrial market - elaboration of main exploitation routs, detailed assessment of exploitation risks and contingency plans. Identification of possible markets and defining a business strategy might be carried out by a consulting company specialized in market analysis.
- **12 month after projects start** – Consultations with potential end-users to verify their needs.
- **After completion of the ACSI Hub System iteration I and II –** Collecting technical feedback and verifying the system prototype (refer to WP5 – Validation).
- **18 months after project start** - Detailed business plans, detailed identification of potential markets and competitive environment, continuation of dissemination.
- **Up to project completion:** Testing and measurements of the results and assessment of business benefits; establishment of commercial agreements between partners on the joint commercialization and exploitation after project completion, including a financial plan; intensive dissemination through demonstration of technology show cases.
- **After project completion:** An implementation plan for future development of the ACSI Hub System in accordance to deliverable 6.4 in the Dissemination and exploitation WP.

3.2.3 Contribution to standards

Although it is not part of the planned 3-year research program for ACSI, it is quite possible that research results from the effort would be appropriate for standardization. In such cases, if they arise, the project coordinator in collaboration with the project partners will consider launching a working group to develop a standard or standard extension. Because of the collective experience of the consortium members with standards activities (see below), the consortium would be well positioned to carry such an effort to successful completion.

There are at least four thematic areas where standards opportunities could arise from the ACSI research.

- The outputs of the project, particularly those of WP1, could form the basis for an extension of the Business Process Modeling Notation (BPMN) for artifact-centric modeling of business processes and service choreographies. This extension would be suitable for submission to the Object Management Group (OMG), which is responsible for the standardization of BPMN.
- Because interoperation hubs offer an alternative to orchestration and choreography, it may be appropriate to bring some of the ACSI technology into W3C standardization. We note that three consortium members have been or are currently involved with W3C standardization activities: Diego Calvanese (FUB) is a member of the W3C working group for the standardization of OWL 2, and Giuseppe De Giacomo (Uniroma1) directly participated in the definition of the W3C recommendation for the OWL 2 Profiles. Also, Rick Hull (IBM) was active in the creation of the Semantic Web Services Language (SWSL) Member Submission, which was submitted in 2005.
- The development of the declarative approach to specifying artifact lifecycles will borrow ideas from OMG's Semantics for Business Vocabulary and Rules (SBVR) standard as appropriate. SBVR is also the standard brought forward by EC's Digital Business Ecosystems initiative. More interestingly, preliminary work in this area by Rick Hull (IBM), Fabiana Fournier (IBM), and others at IBM Research indicates that a declarative language for specifying artifact lifecycles could actually provide useful structure for SBVR rule sets (which otherwise might have a spaghetti-like character). The group's work also suggests that artifact schemas provide a useful way to generate and structure the vocabulary portion of SBVR specifications. It therefore might be appropriate to launch a standardization effort around mechanisms for structuring SBVR vocabularies and rule sets. We note that Mark Linehan (in the group of Rick Hull (IBM)) is currently co-leading a task force that is adding Date and Time to the SBVR standard. Also, since recently, Collibra is an active member of this SBVR committee. Collibra's current software and methodology supports (i) defining business context of important assets in SBVR; (ii) deriving technical models for semantic applications from this business semantics in SBVR; and (iii) generating data transformation and validation services between heterogeneous data sources based on these technical models. Via ACSI, Collibra envisions to expand the application business semantics to three mayor domains: MDM, BAM, and SOA.
- More and more people, both in industry and academia, consider process mining as one of the most important innovations in the field of business process management. It joins ideas of process modeling and analysis on the one hand and data mining and machine learning on the other. Therefore, the IEEE has established a Task Force on Process Mining, cf. http://www.win.tue.nl/ieeetfpm. This Task Force is established in the context of the Data Mining Technical Committee (DMTC) of the Computational Intelligence Society (CIS) of the Institute of Electrical and Electronic Engineers, Inc. (IEEE). The goal of this Task Force is to promote the research, development, education and understanding of process mining. One of the important activities will be to standardize the log format. Based on experiences with MXML (Mining XML format) and SA-MXML (Semantically Annotated MXML) the IEEE Task Force on Process Mining aims to establish a new standard. Here it would be good if ideas

related to artifact centric process mining can be incorporated. W. M. P van der Aalst (TU/e) chairs the IEEE Task Force on Process Mining.

3.2.4 Open Source Software

One of the most effective routes for exploitation of ACSI's results is via open source software licensing. The ACSI project will relate to open source software release as part of its exploitation plan (Deliverable 6.4) aiming to start at an early stage in the project (month 6).

Regarding open source software, we shall consider the following aspects in the exploitation plan:

- We will strive to release as early as possible and incrementally usable material. We will also consider the inclusion of references to/description of their software (residing in a forge) in indexing and reference sites (like Freshmeat.net, Wikipedia or other reference sites).
- In order to increase the chances for adoption of the project technology, we should try to engage a user community as early as possible.

Regarding ACSI's plan for open source software, we anticipate that portions of the ACSI System Hub will be released as open source software. Specifically, the following components will be released as open source software (refer to Figure 1-6 Conceptual architecture of the ACSI Hub System):

- Portions of the Event, Execution, Data access, and Task execution Management components in the Interoperation Hub Engine.
- The basic functionality of the Artifact schema design and Service linkage components in the Design Workbench.

In addition, the components of the Semantic layer, Verification, and Observation-based tools will be released as open source software.

3.2.5 Management of intellectual property

The ACSI project consortium has agreed in principle to manage the knowledge and intellectual property rights as indicated in the following subsections.

The project aims to enrich and expand the state-of-art and state-of-practice of reliable software design and deployment methods for service interoperation in Europe. In order to carry out the work, the partners will develop and share know-how and technologies in many forms, including, but not limited to algorithms, tools, experiences and methodologies. The know-how exchanged between the partners may include, in certain cases, background knowledge.

The partners are in agreement on the principles for the management of intellectual property, as summarized herein.

The consortium agreement will be written in such a way that it is possible for all partners to carry out their project work whenever it is dependent on transfer of knowledge from other partners, whether this is Foreground or Background knowledge.

The consortium agreement will protect the legitimate IP interests of all partners by explicitly limiting the rights to Background knowledge for Use after completion of the project and, where required, even limiting the rights to Foreground knowledge developed during the course of the project when there is no need-to-know or need-to-use. However, Partners that do introduce Background during the project will commit to provide Access Rights for a limited period also after completion of the project.

During the period of the project, Access Rights to Foreground for all Partners for Execution of the Project will be on a Royalty Free basis. Access Rights to Background, as needed by the Partners for Execution of the Project, will also be available on a Royalty Free basis.

In this specific project, it is anticipated that portions of the ACSI Hub System will be released as Open Source. As such, it is the intention of the Consortium that the Access Rights to the Foreground which will be released as Open Source for Use after completion of the project will be Royalty Free or on fair and reasonable conditions. Further, it is the intention of the project that the interoperation hub framework, will be deployable as Software as a Service (SaaS).

Tools, methodology documents, benchmarks and case studies will be available to all. All the partners intend to pursue publications of the underlying principles of the technologies embodied in their tools in the appropriate academic conferences. Finally, all knowledge will be managed in accordance with the FP7 Grant Agreement and a consortium agreement that will be timely prepared and signed by all consortium members. The Consortium Agreement will be based on the EICTA model.

Partner References

IBM

[ADH+09] B. Arlien, D. Dams, R. Hull, J. Letourneau, and K. Namjoshi. Telco meets the Web: Programming shared experience services. To appear, Bell Labs Technical Journal, 2009.

[AFA08] A. Bercovici, F. Fournier, and A. J. Wecker, From Business Architecture to SOA Realization using MDD, in European Conference on Model Driven Architecture Foundations and Applications (ECMDA2008), pp 382-393, 2008.

[AFF+08] A. Bercovici, A. Fisher, F. Fournier, G. Rackham, N. Razinkov, and I. Skarbovsky, A Method for Service Center Architecture Based on Industry Standards, in IEEE International Conference on Services Computing (SCC 2008), 433-440, 2008.

[AH87] S. Abiteboul and R. Hull. IFO: A formal semantic database model. ACM Trans. on Database Systems, 12(4):525–565, Dec. 1987.

[AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu. Foundations of Databases. Addison Wesley Publ. Co., 1995.

[BCDG+05] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Massimo Mecella. Automatic composition of transition-based Semantic Web services with messaging. In Proc. of the 31st Int. Conf. on Very Large Data Bases (VLDB 2005), pages 613–624, 2005.

[BGH+07] K. Bhattacharya, C. Gerede, R. Hull, R. Liu, and J. Su. Towards formal analysis of artifactcentric business process models. In Proc. of 5th Int. Conference on Business Process Management (BPM 2007), 2007.

[BHS09] K. Bhattacharya, R. Hull, and J. Su. A Data-centric Design Methodology for Business Processes. In J. Cardoso and W.M.P. van der Aalst, editors, Handbook of Research on Business Process Management. Information Science Publishing, an imprint of IGI Global, Hershey, PA, USA, 2009.

[CDGHS09] Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Jianwen Su. Artifact-centric workflow dominance. In Proc. of the 7th Int. Joint Conf. on Service Oriented Computing (ICSOC 2009), volume 5900 of Lecture Notes in Computer Science, pages 130–143. Springer, 2009.

[DHH+07] R. Dinoff, T. Ho, R. Hull, B. Kumar, D. Lieuwen, P. Santos, and H. Ren. Intuitive network applications: Learning for personalized converged services involving social networks. J. of Computers, 2(6), 2007.

[DHPV09] Alin Deutsch, Richard Hull, Fabio Patrizi, and Victor Vianu. Automatic verification of datacentric business processes. In Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009), pages 252–267, 2009.

[FFG+08] A. Fisher, F. Fournier, D. Gilat, G. Rackham, N. Razinkov, and S. Wasserkrug, A top-down approach to aligning service centers with business processes.

International Journal of Services Operations and Informatics (IJSOI), 3(3/4):207-220, 2008.

[FHS09] Christian Fritz, Richard Hull, and Jianwen Su. Automatic construction of simple artifact-based business processes. In Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009), pages 225–238, 2009.

[GHJ96] S. Ghandeharizadeh, R. Hull, and D. Jacobs. Heraclitus: Elevating deltas to be first-class citizens in a database programming language. ACM Trans. on Database Systems, 21(3):370–426, 1996.

[HK87] R. Hull and R. King. Semantic database modeling: Survey, applications, and research issues. ACM Computing Surveys, 19(3):201–260, September 1987.

[HKL+05] R. Hull, B. Kumar, D. Lieuwen, P. Patel-Schneider, A. Sahuguet, S. Varadarajan, and A. Vyas. Improving user experience through rule-based service customization. Intl. J. of Cooperative Information Systems, 14(4), December 2005.

[HLS+99] R. Hull, F. Llirbat, E. Simon, J. Su, G. Dong, B. Kumar, and G. Zhou. Declarative workflows that support easy modification and dynamic browsing. In Proc. Int. Joint Conf. on Work Activities Coordination and Collaboration, 1999.

[HNN09] Richard Hull, Nanjangud Narendra, and Anil Nigam. Facilitating workflow interoperation using artifact-centric hubs. In Proc. of the 7th Int. Joint Conf. on Service Oriented Computing (ICSOC 2009), volume 5900 of Lecture Notes in Computer Science. Springer, 2009.

[HS95] R. Hull and J. Su. Domain independence and the relational calculus. Acta Informatica, 31(6):513–524, September 1995.

[Hul86] R. Hull. Relative information capacity of simple relational schemata. SIAM Journal of Computing, 15(3):856–886, August 1986.

[Hul08] Richard Hull. Artifact-centric business process models: Brief survey of research results and challenges. In Proc. of the 7th Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE 2008), volume 5332 of Lecture Notes in Computer Science, pages 1152–1163. Springer,2008.

[HY84] R. Hull and C. K. Yap. The Format model: A theory of database organization. Journal of the ACM, 31(3):518–537, 1984.

Università Degli Studi Di Roma La Sapienza (Uniroma1)

[BCGP08] Daniela Berardi, Fahima Cheikh, Giuseppe De Giacomo, Fabio Patrizi: Automatic Service Composition via Simulation. Int. J. Found. Comput. Sci. 19(2): 429-451 (2008)

[CGHS09] Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Jianwen Su. Artifact-Centric Workflow Dominance. ICSOC 2009.

[CGL+08a] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati: Path-Based Identification Constraints in Description Logics. KR 2008: 231-241

[CGL+08b] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati: Inconsistency tolerance in P2P data integration: An epistemic logic approach. Inf. Syst. 33(4-5): 360-384 (2008)

[CGLR08] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Riccardo Rosati: View-Based Query Answering over Description Logic Ontologies. KR 2008: 242-251

[CGLV09] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Moshe Y. Vardi: An Automata-Theoretic Approach to Regular XPath. DBPL 2009: 18-35

[CGV05] Diego Calvanese, Giuseppe De Giacomo, Moshe Y. Vardi: Decidable containment of recursive queries. Theor. Comput. Sci. 336(1): 33-56 (2005)

[GLL00] Giuseppe De Giacomo, Yves Lesperance, Hector J. Levesque: ConGolog, a concurrent programming language based on the situation calculus. Artif. Intell. 121(1-2): 109-169 (2000)

[GLLR07] Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati: On reconciling data exchange, data integration, and peer data management. PODS 2007: 133-142

[GLPR07] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, Riccardo Rosati: On the Approximation of Instance Level Update and Erasure in Description Logics. AAAI 2007: 403-408

[GLPR09] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, Riccardo Rosati. On Instance-Level Update and Erasure in Description Logic Ontologies. J.Logic and Computation, 19(5): 745-770, 2009

[GMP09] Giuseppe De Giacomo, Riccardo De Masellis, Fabio Patrizi. Composition of Partially Observable Services Exporting their Behaviour. ICAPS 2009.

[GS07] Giuseppe De Giacomo, Sebastian Sardina: Automatic Synthesis of New Behaviors from a Library of Available Behaviors. IJCAI 2007: 1866-1871

[L02] Maurizio Lenzerini: Data Integration: A Theoretical Perspective. PODS 2002: 233-246

[LGLM09] Massimiliano de Leoni, Giuseppe De Giacomo, Yves Lesperance, Massimo Mecella: On-line adaptation of sequential mobile processes running concurrently. SAC 2009: 1345-1352

[LGO08] Yves Lesperance, Giuseppe De Giacomo, Atalay Nafi Ozgovde: A model of contingent planning for agent programming languages. AAMAS (1) 2008: 477-484

[LMG07] Massimiliano de Leoni, Massimo Mecella, Giuseppe De Giacomo: Highly Dynamic Adaptation in Process Management Systems Through Execution Monitoring. BPM 2007: 182-197

[SG08] Sebastian Sardina, Giuseppe De Giacomo: Realizing Multiple Autonomous Agents through Scheduling of Shared Devices. ICAPS 2008: 304-312

[SG09] Sebastian Sardina, Giuseppe De Giacomo: Composition of ConGolog Programs. IJCAI 2009: 904-910

[SGLL04] Sebastian Sardina, Giuseppe De Giacomo, Yves Lesperance, Hector J. Levesque: On the Semantics of Deliberation in Indigolog - from Theory to Implementation. Ann. Math. Artif. Intell. 41(2-4): 259-299 (2004)

[SGLL06] Sebastian Sardina, Giuseppe De Giacomo, Yves Lesperance, Hector J. Levesque: On the Limits of Planning over Belief States under Strict Uncertainty. KR 2006: 463-471

[SPG08] Sebastian Sardina, Fabio Patrizi, Giuseppe De Giacomo: Behavior Composition in the Presence of Failure. KR 2008: 640-650

Libera Università Di Bolzano (FUB)

[ACK+ 07] Alessandro Artale, Diego Calvanese, Roman Kontchakov, Vladislav Ryzhikov, and Michael Zakharyaschev. Reasoning over extended ER models. In Proc. of the 26th Int. Conf. on Conceptual Modeling (ER 2007), volume 4801 of Lecture Notes in Computer Science, pages 277-292. Springer, 2007.

[ACKZ07] Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyaschev. DL-Lite in the light of first-order logic. In Proc. of the 22nd Nat. Conf. on Artificial Intelligence (AAAI 2007), 2007.

[ACKZ09] Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyaschev. The DL-Lite family and relations. J. of Artificial Intelligence Research, 36:1-69, 2009.

[AF09] Alessandro Artale and Enrico Franconi. Foundations of temporal conceptual data models. In A. Borgida, V. Chaudhri, P. Giorgini, and E. Yu, editors, Conceptual Modeling: Foundations and Applications, volume 5600 of Lecture Notes in Computer Science, pages 10-35. Springer, 2009.

[AGK08] Alessandro Artale, Nicola Guarino, and Maria Keet. Formalising temporal constraints on part-whole relations. In Gerhard Brewka and Jerome Lang, editors, Proc. of the 11th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2008), pages 3-14, 2008.

[AL04] A. Artale and C. Lutz. A correspondence between temporal description logics. J. of Applied Non-Classical Logics, 14(1-2):209-233, 2004. Lavoisier.

[ALT07] Alessandro Artale, Carsten Lutz, and David Toman. A description logic of change. In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), 2007.

[APS07] Alessandro Artale, Christine Parent, and Stefano Spaccapietra. Evolving objects in temporal information systems. Ann. of Mathematics and Artificial Intelligence, 50:5-38, 2007. Springer.

[Art06] Alessandro Artale. Reasoning on temporal class diagrams: Undecidability results. Ann. of Mathematics and Artificial Intelligence, 46:265-288, 2006. Springer.

[BCDG+ 04] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Massimo Mecella. Synthesis of underspecified composite e-Services based on automated reasoning. In Proc. of the 2nd Int. Conf. on Service Oriented Computing (ICSOC 2004), pages 105-114, 2004.

[BCDG05a] Daniela Berardi, Diego Calvanese, and Giuseppe De Giacomo. Reasoning on UML class diagrams. Artificial Intelligence, 168(1-2):70-118, 2005.

[BCDG+ 05b] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Massimo Mecella. Automatic composition of transition-based semantic web services with messaging. In Proc. of the 31st Int. Conf. on Very Large Data Bases (VLDB 2005), pages 613-624, 2005.

[BCDG+ 05c] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Massimo Mecella. Automatic service composition based on behavioral descriptions. Int. J. of Cooperative Information Systems, 14(4):333-376, 2005.

[BCDGM05] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, and Massimo Mecella. Composition of services with nondeterministic observable behavior. In Proc. of the 3rd Int. Conf. on Service Oriented Computing (ICSOC 2005), pages 520-526, 2005.

[BCM+ 07] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, 2nd edition, 2007.

[CDGL+ 06] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. In Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006), pages 260-270, 2006.

[CDGL+ 07a] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. EQL-Lite: Effective first-order query processing in description logics. In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), pages 274-279, 2007.

[CDGL+ 07b] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. J. of Automated Reasoning, 39(3):385-429, 2007.

[CDGL+08a] Diego Calvanese, Giuseppe De Giacomo, and Maurizio Lenzerini. Conjunctive query containment and answering under description logics constraints. ACM Trans. on Computational Logic, 9(3):22.1-22.31, 2008.

[CDGL+08b] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Massimo Mecella, and Fabio Patrizi. Automatic service composition and synthesis: the Roman Model. Bull. of the IEEE Computer Society Technical Committee on Data Engineering, 31(3):18-22, 2008.

[CDGLV07] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. View-based query processing: On the relationship between rewriting, answering and losslessness. Theoretical Computer Science, 371(3):169-182, 2007.

[KA08] Maria Keet and Alessandro Artale. Representing and reasoning over a taxonomy of part-whole relations. Applied Ontology, 3(1-2):91-110, 2008. IOS Press.

[PLC+ 08] Antonella Poggi, Domenico Lembo, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. Linking data to ontologies. J. on Data Semantics, X:133-173, 2008.

Imperial College of Science, Technology and Medicine (Imperial)

[BL09] F. Belardinelli, A. Lomuscio. Quantified Epistemic Logic for reasoning about knowledge in Multi-Agent Systems. Artificial Intelligence. Vol 173(9-10), pp 982-1013. Elsevier.

[BL09] F. Belardinelli, A. Lomuscio. First-Order Linear-time Epistemic Logic Group Knowledge: An Axiomatisation of the Monodic Fragment. Proceedings of the 16th Workshop on Logic, Language, Information and Computation. Tokyo, Japan. Lecture Notes in Computer Science. Vol 5514 pp 140-154. Springer.

[BL08] F. Belardinelli, A. Lomuscio. A complete quantified temporal epistemic logic. Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR08). Sydney, Australia. pp 705-714. AAAI Press.

[CDLQ09] M. Cohen, M. Dam, A. Lomuscio, H. Qu. A data symmetry reduction technique for temporal-epistemic logic. Proceedings of the 7th International Symposium on Automated Technology for Verification and Analysis (ATVA09). Macao, SAR. China. Lecture Notes in Computer Science. To appear. Springer.

[CDLR09] M. Cohen, M. Dam, A. Lomuscio, F. Russo Abstraction in model checking multi-agent systems . Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS09). Budapest, Hungary. IFAAMAS Press.

[EL09a] J. Ezekiel, A. Lomuscio. An automated approach to verifying diagnosability in multi-agent systems. Proceedings of the 7th IEEE International Conference on Software Engineering and Formal Methods (SEFM09). Hanoi, Vietnam. IEEE Press. To appear.

[EL09b] J. Ezekiel, A. Lomuscio. Combining fault injection and model checking to verify fault tolerance in multi-agent systems. Proceedings of the 8th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS09). Budapest, Hungary. IFAAMAS Press.

[HLW06] W. van der Hoek, A. Lomuscio, M. Wooldridge. On the complexity of practical ATL model checking knowledge, strategies, and games in multi-agent systems. Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS06). Hakodate, Japan. ACM Press.

[KLP04] M. Kacprzak, A. Lomuscio, W. Penczek. From bounded to unbounded model checking for interpreted systems. Fundamenta Informaticae. Volume 63 (2,3), pp 107-308. 2004.

[LPR07] A. Lomuscio, C. Pecheur, F. Raimondi. Verification of knowledge and time with NuSMV. Proceedings of the 20th International Conference on Artificial Intelligence (IJCAI07), Hyderabadad, India.

[LPW07] A. Lomuscio, W. Penczek, B. Wozna. Bounded model checking knowledge and real time. Artificial Intelligence. Vol 171(16-17), pp 1011-1038. Elsevier.

[LQR09] A. Lomuscio, H. Qu, F. Raimondi MCMAS: A model checker for the verification of multi-agent systems. Proceedings of the 21th International Conference on Computer Aided Verification (CAV 2009). Grenoble, France. Lecture Notes in Computer Science. Vol 5643 pp 682-688. Springer.

[LQS08a] A. Lomuscio, H. Qu, M. Solanki. Towards verifying contract regulated service composition. Proceedings of the 8th International Conference on Web Services (ICWS08). Beijing, China. pp 255-261. IEEE Press.

[LQS08b] A. Lomuscio, H. Qu, M. Solanki. Towards verifying compliance in agent-based web service compositions. Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS08). Estoril, Portugal. IFAAMAS Press.

[LQSS07] A. Lomuscio, H. Qu, M. Sergot, M. Solanki. Verifying Temporal Epistemic properties of Web service compositions. Proceedings of the 5th International Conference on Service Oriented Computing (ICSOC07). Vienna, Austria. Springer LNCS Vol 4749, pp 456-461.

[LR06a] A. Lomuscio, F. Raimondi. The complexity of model checking concurrent programs against CTLK specifications. Post-proceedings of the 4th workshop on Declarative Agent Languages and Technologies (DALT06), Springer LNCS Vol 4327, pp 29-42.

[LR06b] A. Lomuscio, F. Raimondi. Model checking knowledge, strategies, and games in multi-agent systems. Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent systems (AAMAS06). Hakodate, Japan. ACM Press.

[LS03] A. Lomuscio, M. Sergot. Deontic Interpreted Systems. Studia Logica (Special Issue on The Dynamics of Knowledge). W. van der Hoek and M. Wooldridge eds. Kluwer Academic Publishers. Volume 75. 2003.

[LS04] A. Lomuscio, M. Sergot. A formulation of violation, error recovery, and enforcement in the bit transmission problem. Journal of Applied Logic (Selected articles from DEON02 - London). A.J.I. Jones, J. Horty eds. Elsevier. Volume 1. (A previous version of this paper appeared in the proceedings of DEON02). 2004.

[LS09] A. Lomuscio, M. Solanki. Mapping OWL-S processes to multi agent systems: A verification oriented approach. Proceedings of the 6th European Semantic Web Conference (ESWC 2009). Lecture notes in Computer Science. Vol 5554 pp 578-592. Springer. (A previous version of this paper appeared in the Proceedings of the 4th International IEEE Workshop on Service Oriented Architectures in Converging Networked Environments (IEEE Press)).

Technische Universiteit Eindhoven (TUE)

[A98] W.M.P. van der Aalst. The Application of Petri Nets to Workflow Management. The Journal of Circuits, Systems and Computers, 8(1):21-66, 1998.

[A03] W.M.P. van der Aalst. Don't go with the flow: Web services composition standards exposed. IEEE Intelligent Systems, 18(1):72-76, 2003.

[ABEW00] W.M.P. van der Aalst, P. Barthelmess, C.A. Ellis, and J. Wainer. Workflow Modeling using Proclets. In O. Etzion and P. Scheuermann, editors, 7th International Conference on Cooperative Information Systems (CoopIS 2000), volume 1901 of Lecture Notes in Computer Science, pages 198-209. Springer-Verlag, Berlin, 2000.

[ABEW01] W.M.P. van der Aalst, P. Barthelmess, C.A. Ellis, and J. Wainer. Proclets: A Framework for Lightweight Interacting Workflow Processes. International Journal of Cooperative Information Systems, 10(4):443-482, 2001.

[AB02] W.M.P. van der Aalst and T. Basten. Inheritance of Workflows: An Approach to Tackling Problems Related to Change. Theoretical Computer Science, 270(1-2):125-203, 2002.

[ABD05] W.M.P. van der Aalst, H.T. de Beer, and B.F. van Dongen. Process Mining and Verification of Properties: An Approach based on Temporal Logic. In R. Meersman and Z. Tari et al., editors, On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE: OTM Confederated International Conferences, CoopIS, DOA, and ODBASE 2005, volume 3760 of Lecture Notes in Computer Science, pages 130-147. Springer-Verlag, Berlin, 2005.

[ADO+08] W.M.P. van der Aalst, M. Dumas, C. Ouyang, A. Rozinat, and H.M.W. Verbeek. Conformance Checking of Service Behavior. ACM Transactions on Internet Technology, 8(3):29-59, 2008.

[ADH+03] W.M.P. van der Aalst, B.F. van Dongen, J. Herbst, L. Maruster, G. Schimm, and A.J.M.M. Weijters. Workflow Mining: A Survey of Issues and Approaches. Data and Knowledge Engineering, 47(2):237-267, 2003.

[AHEA06] M. Adams, A.H.M. ter Hofstede, D. Edmond, and W.M.P. van der Aalst. Worklets: A Service-Oriented Implementation of Dynamic Flexibility in Workflows. In R. Meersman and Z. Tari et al., editors, On the Move to Meaningful Internet Systems 2006, OTM Confederated International Conferences, 14th International Conference on Cooperative Information Systems (CoopIS 2006), volume 4275 of Lecture Notes in Computer Science, pages 291-308. Springer-Verlag, Berlin, 2006.

[AH05] W.M.P. van der Aalst and A.H.M. ter Hofstede. YAWL: Yet Another Workflow Language. Information Systems, 30(4):245-275, 2005.

[AHKB03] W.M.P. van der Aalst, A.H.M. ter Hofstede, B. Kiepuszewski, and A.P. Barros. Workflow Patterns. Distributed and Parallel Databases, 14(1):5-51, 2003.

[AKV03] W.M.P. van der Aalst, A. Kumar, and H.M.W. Verbeek. Organizational Modeling in UML and XML in the context of Workflow Systems. In H. Haddad and G. Papadopoulos, editors, Proceedings of the 18th Annual ACM

Symposium on Applied Computing (SAC 2003), pages 603-608. ACM Press, 2003.

[AMR09] W.M.P. van der Aalst, R.S. Mans, and N.C. Russell. Workflow Support Using Proclets: Divide, Interact, and Conquer. IEEE Bulletin of the Technical Committee on Data Engineering, 32(3):16-22, 2009.

[AW04] W.M.P. van der Aalst and A.J.M.M. Weijters. Process Mining: A Research Agenda. Computers in Industry, 53(3):231-244, 2004.

[AWM04] W.M.P. van der Aalst, A.J.M.M. Weijters, and L. Maruster. Workflow Mining: Discovering Process Models from Event Logs. IEEE Transactions on Knowledge and Data Engineering, 16(9):1128-1142, 2004.

[DMV+05] B.F. van Dongen, A.K. Alves de Medeiros, H.M.W. Verbeek, A.J.M.M. Weijters, and W.M.P. van der Aalst. The ProM framework: A New Era in Process Mining Tool Support. In G. Ciardo and P. Darondeau, editors, Application and Theory of Petri Nets 2005, volume 3536 of Lecture Notes in Computer Science, pages 444-454. Springer-Verlag, Berlin, 2005.

[DAH05] M. Dumas, W.M.P. van der Aalst, and A.H.M. ter Hofstede. Process-Aware Information Systems: Bridging People and Software through Process Technology. Wiley & Sons, 2005.

[GA07] C.W. Günther and W.M.P. van der Aalst. Fuzzy Mining: Adaptive Process Simplification Based on Multi-perspective Metrics. In G. Alonso, P. Dadam, and M. Rosemann, editors, International Conference on Business Process Management (BPM 2007), volume 4714 of Lecture Notes in Computer Science, pages 328-343. Springer-Verlag, Berlin, 2007.

[MAW08] [A.K. Alves de Medeiros, W.M.P. van der Aalst, and A.J.M.M. Weijters. Quantifying Process Equivalence Based on Observed Behavior. Data and Knowledge Engineering, 64(1):55-74, 2008.

[MWA07] A.K. Alves de Medeiros, A.J.M.M. Weijters, and W.M.P. van der Aalst. Genetic Process Mining: An Experimental Evaluation. Data Mining and Knowledge Discovery, 14(2):245-304, 2007.

[MNA07] J. Mendling, G. Neumann, and W.M.P. van der Aalst. Understanding the Occurrence of Errors in Process Models Based on Metrics. In F. Curbera, F. Leymann, and M. Weske, editors, Proceedings of the OTM Conference on Cooperative information Systems (CoopIS 2007), volume 4803 of Lecture Notes in Computer Science, pages 113-130. Springer-Verlag, Berlin, 2007.

[RA07] M. Rosemann and W.M.P. van der Aalst. A Configurable Reference Modeling Language. Information Systems, 32(1):1-23, 2007.

[RMSA09] A. Rozinat, R.S. Mans, M. Song, and W.M.P. van der Aalst. Discovering Simulation Models. Information Systems, 34(3):305-327, 2009.

[SAA99] A.P. Sheth, W.M.P. van der Aalst, and I.B. Arpinar. Processes Driving the Networked Economy: ProcessPortals, ProcessVortex, and Dynamically Trading Processes. IEEE Concurrency, 7(3):18-31, 1999.

[SWDA08] H. Schonenberg, B. Weber, B.F. van Dongen, and W.M.P. van der Aalst. Supporting Flexible Processes Through Recommendations Based on History.

In M. Dumas, M. Reichert, and M.C. Shan, editors, International Conference on Business Process Management (BPM 2008), volume 5240 of Lecture Notes in Computer Science, pages 51-66. Springer-Verlag, Berlin, 2008.

Tartu Ulikool (UT)

[ADO08] W.M.P. van der Aalst, M. Dumas, C. Ouyang, A. Rozinat, and H.M.W. Verbeek. “Conformance Checking of Service Behavior.” ACM Transactions on Internet Technology 8(3), May 2008.

[ARD07] W.M.P. van der Aalst, M. Rosemann, and M. Dumas. “Deadline-based Escalation in Process-Aware Information Systems” Decision Support Systems 43(2):492-511, March 2007. Elsevier Science Publications

[BDS05] B. Benatallah, M. Dumas, and Q.Z. Sheng: “Facilitating the Rapid Development and Scalable Orchestration of Composite Web Services.” Distributed and Parallel Databases, 15(1):5-37, January 2005. Kluwer Academic Publishers.

[DBM08] M. Dumas, B. Benatallah and H.R. Motahari-Nezhad. Web Service Protocols: Compatibility and Adaptation. IEEE Data Engineering Bulletin 31(3):40-44, September 2008.

[DD04] R. Dijkman and M. Dumas: “Service-oriented Design: A Multi-viewpoint Approach.” International Journal of Cooperative Information Systems 13(4):337-368, December 2004. World Scientific.

[OAD+10] C. Ouyang, W.M.P. van der Aalst, M. Dumas, A.H.M. ter Hofstede, and J. Mendling. “From Business Process Models to Process-Oriented Software Systems”. ACM Transactions on Software Engineering Methodology 19(10), August 2010.

[ORH+08] C. Ouyang, M. La Rosa, A. H. M. ter Hofstede, M. Dumas, K. Shortland. “Towards Web-Scale Workflows for Film Production.” IEEE Internet Computing 12(5):53-61, September 2008.

[OVA+07] C. Ouyang, H.M.W. Verbeek, W.M.P. van der Aalst, S. Breutel, M. Dumas, A.H.M. ter Hofstede. Formal Semantics and Analysis of Control Flow in WS-BPEL. Science of Computer Programming 67(2-3):125-332, July 2007. Elsevier Science Publications.

[RDHI08] G. Redding, M. Dumas, A.H.M. ter Hofstede, and A. Iordachescu. Generating Business Process Models from Object Behavior Models. Information Systems Management 25(4):3-15. October 2008. Taylor & Francis.

[RKM04] J. Rao, P. Küngas, M. Matskin: Logic-based Web Services Composition: From Service Description to Process Model. International Conference on Web Services (ICWS) 2004, pp. 446-453

[RKM06] J. Rao, P. Küngas, M. Matskin: Composition of Semantic Web services using Linear Logic theorem proving. Information Systems 31(4-5):340-360, 2006

[ZBN+04] L. Zeng, B. Benatallah, A. H.H. Ngu, M. Dumas, J. Kalagnanam, and H. Chang: “QoS-aware Middleware for Web Services Composition.” IEEE Transactions on Software Engineering 30(5):311-327, May 2004. IEEE Computer Society.

Collibra NV/SA (Collibra)

[BMST07] Dang Bui Bach, Robert Meersman, Peter Spyns, Damien Trog: Mapping OWL-DL into ORM/RIDL. OTM Workshops (1) 2007: 742-751

[CLMM08] S. Christiaens, P. De Leenheer, A. de Moor, and R. Meersman. Business use case: Ontologising competencies in an interorganisational setting. In M. Hepp, P. De Leenheer, A. de Moor, and Y. Sure, editors, Ontology Management for the Semantic Web, Semantic Web Services, and Business Applications, from Semantic Web and Beyond: Computing for Human Experience. Springer, 2008.

[GSCJ09] G. Van Grootel, P. Spyns, S. Christiaens, B. Jorg. Business Semantics Management supports government innovation information portal. In Proc. Of OTM-ORM 2009. Forthcoming.

[HLMY08] Hepp, M., De Leenheer, P., de Moor, A., and Sure, Y. (eds.) Ontology Management for the Semantic Web, Semantic Web Services, and Business Applications. Springer. ISBN: 978-0-387-69899-1, 2008.

[LC08] P. De Leenheer and S. Christiaens. Challenges and opportunities for more meaningful and sustainable internet systems. In Proc. of the International Future Internet Symposium, LNCS, Vienna, Austria, 2008. Springer.

[LCM09] P. De Leenheer, S. Christiaens, and R. Meersman. Business semantics management: a case study for competency-centric HRM. Journal of Computers For Industry, forthcoming, 2009.

[LD09] De Leenheer, P. and Debruyne, C. (2009) Towards Social Performance Indicators for Community-based Ontology Evolution. In Proc. of ISWC Workshop on Collaborative Construction, Management and Linking of Structured Knowledge (CK2008), CEUR, forthcoming

[Lee08] De Leenheer, P. (2008) Keynote: Ontological Foundations for Evolving Agent Communities. In Proc of 32nd Annual IEEE International Computer Software and Applications Conference (COMPSAC 2008), IEEE Press, pp. 523-528

[Lee09a] P. De Leenheer. On Community-based Ontology Evolution. PhD thesis, Vrije Universiteit Brussel, Brussels, Belgium, 2009.

[Lee09b] De Leenheer, P. (2009) Ontology Elicitation. In Encyclopedia of Database Systems, eds. Liu, L. and Ôzsu, T., Springer, forthcoming Spring 2009.

[LM05] De Leenheer, P. and de Moor, A. (2005) Context-driven Disambiguation in Ontology Elicitation. In Shvaiko P. & Euzenat J.,(eds.), Context and Ontologies: Theory, Practice and Applications, AAAI Technical Report WS-05-01 (Pittsburgh, Pennsylvania), AAAI Press, pp. 17–24.

[LM07a] P. De Leenheer and T. Mens. Using graph transformation to support collaborative ontology evolution. In A. Schürr, M. Nagl, and A. Zündorf, editors, Proc. of Agtive (Kassel, Germany), volume 5088 of LNCS, pages 44–58. Springer, 2007.

[LM07b] De Leenheer, P. and Meersman, R. (2007) Towards Community-driven Evolution of Knowledge-intensive Systems. In Proc. of the 6th Int'l Conf. on

Ontologies, DataBases, and Applications of Semantics (ODBASE 2007) (Vilamoura, Portugal), LNCS, Springer

[LM08] P. De Leenheer and T. Mens. Ontology evolution: State of the art and future directions. In M. Hepp, P. De Leenheer, A. de Moor, and Y. Sure, editors, Ontology Management for the Semantic Web, Semantic Web Services, and Business Applications. Springer, 2008.

[LMM07] P. De Leenheer, A. de Moor, and R. Meersman. Context dependency management in ontology engineering: a formal approach. LNCS Journal on Data Semantics, 8:26–56, 2007.

[Mee99] R. Meersman. The use of lexicons and other computer-linguistic tools in semantics, design and cooperation of database systems. In Proc.of the Conf. on Cooperative Database Systems (CODAS99), pages 1–14. Springer, 1999.

[MLM06] A. de Moor, P. De Leenheer, and R. Meersman. DOGMA-MESS: A meaning evolution support system for interorganizational ontology engineering. In In Proc. of the 14th Int'l Conference on Conceptual Structures (ICCS 2006) (Aalborg, Denmark), LNAI 4068, pages 189–203. Springer Verlag, 2006.

[SMJ02] P. Spyns, R. Meersman, and M. Jarrar. Data modeling versus ontology engineering. SIGMOD Record, 31(4):12–17, 2002.

[TTM07] Damien Trog, Yan Tang, Robert Meersman: Towards Ontological Commitments with Omega -RIDL Markup Language. RuleML 2007: 92-106

[TVC+06] Trog, D., Vereecken, J., Christiaens, S., De Leenheer, P., and Meersman, R. (2006) T-Lex: a Role-based Ontology Engineering Tool. In Proc. of the On The Move to Meaningful Internet Systems Workshops (OTM2006) (Montpelier, France) , LNCS 4278, Springer, pp. 1191-1200.

References

[AA+99] P. Abdulla, A. Annichini, S. Bensalem, A. Bouajjani, P. Habermehl, and Y. Lakhnech. Verification of infinite-state systems by combining abstraction and reachability analysis. In Proc. of the 11th Int. Conf. on Computer Aided Verification (CAV'99), pages 146-159, 1999.

[ABGM09] S. Abiteboul, P. Bourhis, A. Galland, and B. Marinoiu. The AXML artifact model. In Proc. of the 16th Int. Symp. on Temporal Representation and Reasoning (TIME 2009), 2009.

[ACKM04] G. Alonso, F. Casati, H. Kuno, and V. Machiraju. Web Services, Concepts Architectures and Applications.

[AGL98] Rakesh Agrawal, Dimitrios Gunopulos, and Frank Leymann. Mining process models from workflow logs. In Proc. of the 6th Int. Conf. on Extending Database Technology (EDBT'98), pages 469-483, 1998.

[AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. J. of the ACM, 49(5):672-713, 2002.

[AHV95] Serge Abiteboul, Richard Hull, and Victor Vianu. Foundations of Databases. Addison Wesley Publ. Co., 1995.

[AKRZ09] Alessandro Artale, Roman Kontchakov, Vladislav Ryzhikov, and Michael Zakharyaschev. DL-Lite with temporalised concepts, rigid axioms and roles. In Proc. of the 7th Int. Symp. on Frontiers of Combining Systems (FroCoS 2009), volume 5749 of Lecture Notes in Computer Science, pages 133-148. Springer, 2009.

[AT08] Alessandro Artale and David Toman. Decidable reasoning over timestamped conceptual models. In Proc. of the 21st Int. Workshop on Description Logic (DL 2008), volume 353 of CEUR Electronic Workshop Proceedings, http://ceur-ws.org/, 2008.

[BBL09] Franz Baader, Andreas Bauer, and Marcel Lippmann. Runtime verification using a temporal de- scription logic. In Proc. of the 7th Int. Symp. on Frontiers of Combining Systems (FroCoS 2009), volume 5749 of Lecture Notes in Computer Science, pages 149-164. Springer, 2009.

[BCDG+05] Daniela Berardi, Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Massimo Mecella. Automatic composition of transition-based Semantic Web services with messaging. In Proc. of the 31st Int. Conf. on Very Large Data Bases (VLDB 2005), pages 613-624, 2005.

[BCK+07] K. Bhattacharya, N. S. Caswell, S. Kumaran, A. Nigam, and F. Y. Wu. Artifact-centered operational modeling: Lessons from customer engagements. IBM Systems Journal, 46(4):703-721, 2007.

[BDW07] M. Born, F. D"orr, and I. Weber. User-friendly semantic annotation in business process modeling. In Proc. of the Workshop on Human-friendly Service Description, Discovery, and Matchmaking (Hf-SDDM), 2007. Included in

Workshop Proc. of Int. Conf. on Web Information Systems Engineering (WISE).

[BF08] Tevfik Bultan and Xiang Fu. Choreography Modeling and Analysis with Collaboration Diagrams. Bull. of the IEEE Computer Society Technical Committee on Data Engineering, 31(3):27-30, 2008.

[BGH+07] K. Bhattacharya, C. Gerede, R. Hull, R. Liu, and J. Su. Towards formal analysis of artifact-centric business process models. In Proc. of 5th Int. Conference on Business Process Management (BPM 2007), 2007.

[BGL08] Franz Baader, Silvio Ghilardi, and Carsten Lutz. LTL over description logic axioms. In Proc. of the 11th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2008), pages 684-694, 2008.

[BHS09] K. Bhattacharya, R. Hull, and J. Su. A data-centric design methodology for business processes. In J. Cardoso and W.M.P. van der Aalst, editors, Handbook of Research on Business Process Management. Information Science Publishing, an imprint of IGI Global, Hershey, PA, USA, 2009.

[BLO98] S. Bensalem, Y. Lakhnech, and S. Owre. Computing abstractions of infinite state systems com- positionally and automatically. In Proc. of the 10th Int. Conf. on Computer Aided Verification (CAV'98), pages 319-331, 1998.

[Bry92] Randal E. Bryant. Symbolic Boolean manipulation with ordered binary decision diagrams. ACM Computing Surveys, 24(3):293-318, 1992.

[BSB+08] M. J. Ball, J. S. Silva, S. Bierstock, J. V. Douglas, A. F. Norcio, J. Chakraborty, and J. Srini. Failure to provide clinicians useful IT systems: Opportunities to leapfrog current technologies. Methods of Information in Medicine, 47(1), 2008.

[C+02] NuSMV 2: An OpenSource Tool for Symbolic Model Checking A. Cimatti, E. M. Clark Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani and A. Tacchella. In Proceeding of the 14th International Conference on Computer-Aided Verification (CAV 2002). Copenhagen, Denmark. Lecture Notes in Computer Science. Vol 2404 359-364. Springer. 2002

[C+09] [CCDGL04] Andrea Cali, Diego Calvanese, Giuseppe De Giacomo, and Maurizio Lenzerini. Data integration under integrity constraints. Information Systems, 29:147-163, 2004.

[CDGHS09] Diego Calvanese, Giuseppe De Giacomo, Rick Hull, and Jianwen Su. Artifact-centric workflow dominance. In Proc. of the 7th Int. Joint Conf. on Service Oriented Computing (ICSOC 2009), volume 5900 of Lecture Notes in Computer Science, pages 130-143. Springer, 2009.

[CDGL98] Diego Calvanese, Giuseppe De Giacomo, and Maurizio Lenzerini. On the decidability of query containment under constraints. In Proc. of the 17th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS'98), pages 149-158, 1998.

[CDGL+05] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tailoring OWL for data intensive ontologies. In Proc. of the 1st Int. Workshop on OWL: Experiences and Directions

(OWLED 2005), volume 188 of CEUR Electronic Workshop Proceedings, http://ceur-ws.org/, 2005.

[CDGL+06] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Data complexity of query answering in description logics. In Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006), pages 260-270, 2006.

[CDGL+07a] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. EQL-Lite: Effective first-order query processing in description logics. In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), pages 274-279, 2007.

[CDGL+07b] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The DL-Lite family. J. of Automated Reasoning, 39(3):385-429, 2007.

[CDGL+08a] Diego Calvanese, Giuseppe De Giacomo, and Maurizio Lenzerini. Conjunctive query containment and answering under description logics constraints. ACM Trans. on Computational Logic, 9(3):22.1-22.31, 2008.

[CDGL+08b] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Massimo Mecella, and Fabio Patrizi. Automatic service composition and synthesis: the Roman Model. Bull. of the IEEE Computer Society Technical Committee on Data Engineering, 31(3):18-22, 2008.

[CDGL+09a] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Antonella Poggi, Mariano Rodriguez-Muro, and Riccardo Rosati. Ontologies and databases: The DL-Lite approach. In Sergio Tessaris and Enrico Franconi, editors, Semantic Technologies for Informations Systems - 5th Int. Reasoning Web Summer School (RW 2009), volume 5689 of Lecture Notes in Computer Science, pages 255-356. Springer, 2009.

[CDGL+09b] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Conceptual modeling for data integration. In Alex T. Borgida, Vinay Chaudhri, Paolo Giorgini, and Eric Yu, editors, Conceptual Modeling: Foundations and Applications - Essays in Honor of John Mylopoulos, volume 5600 of Lecture Notes in Computer Science, pages 173-197. Springer, 2009.

[CDGLR07] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. Actions and programs over description logic ontologies. In Proc. of the 20th Int. Workshop on Description Logic (DL 2007), volume 250 of CEUR Electronic Workshop Proceedings, http://ceur-ws. org/, pages 29-40, 2007.

[CDGLR08] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. View-based query answering over description logic ontologies. In Proc. of the 11th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2008), pages 242-251, 2008.

[CDGLV07] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Moshe Y. Vardi. View-based query processing: On the relationship between rewriting, answering and losslessness. Theoretical Computer Science, 371(3):169-182, 2007.

[CGHKS07] Bernardo Cuenca Grau, Ian Horrocks, Yevgeny Kazakov, and Ulrike Sattler. A logical framework for modularity of ontologies. In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), 2007.

[CGP99] Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. Model Checking. The MIT Press, Cambridge, MA, USA, 1999.

[CHP+08] David Cohn, Fenno Heath, Florian Pinel, John Vergo, and Pankaj Dhoolia. Siena: From PowerPoint to Web application in 5 minutes. In Proc. of the 6th Int. Conf. on Service Oriented Computing (ICSOC 2008), volume 5364 of Lecture Notes in Computer Science, pages 722-723. Springer, 2008.

[CL08] Jens Classen and Gerhard Lakemeyer. A logic for non-terminating Golog programs. In Proc. of the 11th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2008), pages 589-599, 2008.

[CL09] Jens Classen and Gerhard Lakemeyer. Tractable first-order Golog with disjunctive knowledge bases. In Proc. of the 9th Int. Symp. on Logical Formalizations of Commonsense Reasoning (Commonsense 2009), 2009.

[CSC08] Javier aCmara, Gwen Sala"un, and Carlos Canal. Composition and run-time adaptation of mismatching behavioural interfaces. J. of Universal Computer Science, 14(13):2182-2211, 2008.

[CW99] Jonathan E. Cook and Alexander L. Wolf. Software process validation: Quantitatively mea- suring the correspondence of a process to a model. ACM Trans. on Software Engineering and Methodology, 8(2):147-176, 1999.

[dB09] Jos de Bruijn. State-of-the-art survey of issues. Deliverable D3.1, OntoRule Project, 2009. Available at http://ontorule-project.eu/.

[DDP99] S. Das, D. Dill, and S. Park. Experience with predicate abstraction. In Proc. of the 11th Int. Conf. on Computer Aided Verification (CAV'99), pages 160-171, 1999.

[DGLL00] Giuseppe De Giacomo, Yves Lesperance, and Hector J. Levesque. ConGolog, a concurrent programming language based on the situation calculus. Artificial Intelligence, 121(1-2):109-169, 2000.

[DGLLR07] Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. On reconciling data exchange, data integration, and peer data management. In Proc. of the 26th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2007), pages 133-142, 2007.

[DGLPR06] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the update of description logic ontologies at the instance level. In Proc. of the 21st Nat. Conf. on Artificial Intelligence (AAAI 2006), pages 1271-1276, 2006.

[DGLPR07] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On the approximation of instance level update and erasure in description logics. In Proc. of the 22nd Nat. Conf. on Artificial Intelligence (AAAI 2007), pages 403-408, 2007.

[DGLPR09] Giuseppe De Giacomo, Maurizio Lenzerini, Antonella Poggi, and Riccardo Rosati. On instance-level update and erasure in description logic ontologies. J.

of Logic and Computation, Special Issue on Ontology Dynamics, 19(5):745-770, 2009.

[DHPV09] Alin Deutsch, Richard Hull, Fabio Patrizi, and Victor Vianu. Automatic verification of data-centric business processes. In Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009), pages 252-267, 2009.

[DMCS05] N. Desai, A. U. Mallya, A. K. Chopra, and M. P. Singh. Processes = Protocols + Policies: A methodology for business process development. In Proc. of the 14th Int. World Wide Web Conf. (WWW 2005), 2005.

[DSV07] Alin Deutsch, Liying Sui, and Victor Vianu. Specification and verification of data-driven web applications. J. of Computer and System Sciences, 73(3):442-474, 2007.

[EC80] E. A. Emerson and E. M. Clarke. Characterizing correctness properties of parallel programs using fixpoints. In Proc. of the 7th Coll. on Automata, Languages and Programming (ICALP'80), pages 169-181, 1980.

[EEO+98] Wolfgang Emmerich, Ernst Ellmer, Birgit Osterholt, Roberto Zicari: Business Object Facilities - A Comparative Analysis. Workshop on Integration of Heterogeneous Software Systems 1998: 1-12

[EF06] Cindy Eisner and Dana Fisman. A Practical Introduction to PSL. Integrated Circuits and Systems. Springer, 2006.

[Eme97] E. Allen Emerson. Model checking and the Mu-calculus. In N. Immerman and P. Kolaitis, editors, Proc. of the DIMACS Symposium on Descriptive Complexity and Finite Model, pages 185-214. American Mathematical Society Press, 1997.

[Eur] Europa. Publications office home page. Available at http://publications.europa.eu/about_us/key_figures/index_en.htm.

[Fag07] Ronald Fagin. Inverting schema mappings. ACM Trans. on Database Systems, 32(4), 2007.

[FHMV95] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. Reasoning about Knowledge. The MIT Press, 1995.

[FHS09] Christian Fritz, Richard Hull, and Jianwen Su. Automatic construction of simple artifact-based business processes. In Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009), pages 225-238, 2009.

[FKP05] Ronald Fagin, Phokion G. Kolaitis, and Lucian Popa. Data exchange: Getting to the core. ACM Trans. on Database Systems, 30(1):174-210, 2005.

[FKPT08] Ronald Fagin, Phokion G. Kolaitis, Lucian Popa, and Wang-Chiew Tan. Quasi-inverses of schema mappings. ACM Trans. on Database Systems, 33(2):1-52, 2008.

[FKPT09] Ronald Fagin, Phokion G. Kolaitis, Lucian Popa, and Wang Chiew Tan. Reverse data exchange: Coping with nulls. In Proc. of the 28th ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2009), pages 23-32, 2009.

[GHLS07] Birte Glimm, Ian Horrocks, Carsten Lutz, and Ulrike Sattler. Conjunctive query answering for the description logic SHIQ . In Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007), pages 399-404, 2007.

[GvdM04] Peter Gammie and Ron van der Meyden. MCK: model checking the logic of knowledge. In Proc. of the 16th Int. Conf. on Computer Aided Verification (CAV 2004), pages 479-483, 2004.

[Hal01] Alon Y. Halevy. Answering queries using views: A survey. Very Large Database J., 10(4):270- 294, 2001.

[Hal03] Alon Y. Halevy. Data integration: A status report. In Proc. of the 10th Conf. on Database Systems for Business, Technology and Web (BTW 2003), pages 24-29, 2003.

[HBHP09] Jorg Hoffmann, Piergiorgio Bertoli, Malte Helmert, and Marco Pistore. Message-based web service composition, integrity constraints, and planning under uncertainty: A new connection. J. of Artificial Intelligence Research, 35:49-117, 2009.

[HKT00] David Harel, Dexter Kozen, and Jerzy Tiuryn. Dynamic Logic. The MIT Press, 2000.

[HMA+08] Stijn Heymans, Li Ma, Darko Anicic, Zhilei Ma, Nathalie Steinmetz, Yue Pan, Jing Mei, Achille Fokoue, Aditya Kalyanpur, Aaron Kershenbaum, Edith Schonberg, Kavitha Srinivas, Cristina Feier, Graham Hench, Branimir Wetzstein, and Uwe Keller. Ontology reasoning with large data repositories. In Martin Hepp, Pieter De Leenheer, Aldo de Moor, and York Sure, editors, Ontology Management, Semantic Web, Semantic Web Services, and Business Applications, volume 7 of Semantic Web And Beyond Computing for Human Experience, pages 89-128. Springer, 2008.

[HNN09] Richard Hull, Nanjangud Narendra, and Anil Nigam. Facilitating workflow interoperation using artifact-centric hubs. In Proc. of the 7th Int. Joint Conf. on Service Oriented Computing (ICSOC 2009), volume 5900 of Lecture Notes in Computer Science. Springer, 2009.

[Hul08] Richard Hull. Artifact-centric business process models: Brief survey of research results and challenges. In Proc. of the 7th Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE 2008), volume 5332 of Lecture Notes in Computer Science, pages 1152-1163. Springer, 2008.

[Kol05] Phokion G. Kolaitis. Schema mappings, data exchange, and metadata management. In Proc. of the 24rd ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2005), pages 61-75, 2005.

[KL08] Oliver Kopp and Frank Leymann. Choreography Design Using WS-BPEL. Bull. of the IEEE Computer Society Technical Committee on Data Engineering, 31(3):31-34, 2008.

[KPS+09] R. Kontchakov, L. Pulina, U. Sattler, T. Schneider, P. Selmer, F. Wolter, and M. Zakharyaschev. Minimal module extraction from DL-Lite ontologies using QBF solvers. In Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009), pages 836-840, 2009.

[Len02] Maurizio Lenzerini. Data integration: A theoretical perspective. In Proc. of the 21st ACM SIGACT SIGMOD SIGART Symp. on Principles of Database Systems (PODS 2002), pages 233-246, 2002.

[Lev84] Hector J. Levesque. Foundations of a functional approach to knowledge representation. Artificial Intelligence, 23:155-212, 1984.

[LLMW06] H. Liu, C. Lutz, M. Milicic, and F. Wolter. Updating description logic ABoxes. In Proc. of the 10th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR 2006), pages 46-56, 2006.

[LM99] M. Leuschel and T. Massart. Infinite state model checking by abstract interpretation and program specialisation. In Selected Papers of the 9th Int. Workshop on Logic Programming Synthesis and Transformation (LOPSTR'99), volume 1817 of Lecture Notes in Computer Science, pages 62-81. Springer, 1999.

[LPSS09] A. Lomuscio, W. Penczek, M. Solanki, and M. Szreter. Runtime monitoring of contract regulated web services. In Proc. of the 12th Int. Workshop on Concurrency, Specification and Programming (CS&P 2009), 2009.

[LQR09] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. MCMAS: A model checker for the verification of multi-agent systems. In Proc. of the 21st Int. Conf. on Computer Aided Verification (CAV 2009), pages 682-688, 2009.

[LRL+97] H. J. Levesque, R. Reiter, Y. Lesperance, F. Lin, and R. Scherl. GOLOG: A logic programming language for dynamic domains. J. of Logic Programming, 31:59-84, 1997.

[LTW09] Carsten Lutz, David Toman, and Frank Wolter. Conjunctive query answering in the description logic el using a relational database system. In Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009), pages 2070-2075, 2009.

[MNBM+07] Hamid R. Motahari Nezhad, Boualem Benatallah, Axel Martens, Francisco Curbera, and Fabio Casati. Semi-automated adaptation of service interactions. In Proc. of the 16th Int. World Wide Web Conf. (WWW 2007), pages 993-1002, 2007.

[MP92] Z. Manna and A. Pnueli. The Temporal Logic of Reactive and Concurrent Systems: Specification. Springer, 1992.

[MPvdA+10] M. Montali, M. Pesic, W.M.P. van der Aalst, F. Chesani, P. Mello, and S. Storari. Declarative Specification and Verification of Service Choreographies. ACM Transactions on the Web, 4(1) 2010.

[NC03] A. Nigam and N. S. Caswell. Business artifacts: An approach to operational specification. IBM Systems Journal, 42(3):428-445, 2003.

[Obj08] Object Management Group (OMG). Semantics of Business Vocabulary and Business Rules (SBVR), version 1.0, January 2008. Available at http://www.omg.org/spec/SBVR/1.0/.

[OCE08] Magdalena Ortiz, Diego Calvanese, and Thomas Eiter. Data complexity of query answering in expressive description logics via tableaux. J. of Automated Reasoning, 41(1):61-98, 2008.

[P+04] T.A. Pardo et al. Modeling the social & technical processes of interorganizational information integration. In Proc. of the 37th Hawaii Int. Conf. on System Sciences, 2004.

[PGGB09] T.A. Pardo, J.R. Gil-Garcia, and G.B. Burke. Information sharing and public health: A case-based look at the ICT expectations-reality gap. In A. Meijer, K. Boersma, and P. Wagenaar, editors, ICTs, Citizens & Governance: After the Hype!, pages 180-197. IOS Press, 2009.

[PL03] Wojciech Penczek and Alessio Lomuscio. Verifying epistemic properties of multi-agent systems via bounded model checking. Fundam. Inform., 55(2):167-185, 2003.

[PLC+08] Antonella Poggi, Domenico Lembo, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. Linking data to ontologies. J. on Data Semantics, X:133-173, 2008.

[PvdA06] M. Pesic and W.M.P. van der Aalst. A Declarative Approach for Flexible Business Processes. In J. Eder and S. Dustdar, editors, Business Process Management Workshops, Workshop on Dynamic Process Management (DPM 2006), volume 4103 of Lecture Notes in Computer Science, pages 169-180. Springer-Verlag, Berlin, 2006.

[RDtHI09] G. Redding, M. Dumas, A.H.M. ter Hofstede, and A. Iordachescu. Modeling flexible processes with business objects. In Proc. of the 11th IEEE Int. Conf. on Commerce and Enterprise Computing (CEC), 2009.

[Rei01] Raymond Reiter. Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems. The MIT Press, 2001.

[RL07] Franco Raimondi and Alessio Lomuscio. Automatic verification of multi-agent systems by model checking via ordered binary decision diagrams. J. Applied Logic, 5(2):235-251, 2007.

[RSE08] Franco Raimondi, James Skene, and Wolfgang Emmerich. Efficient online monitoring of web-service SLAs. In Proc. of the 16th ACM SIGSOFT Int. Symposium on Foundations of Software Engineering (SIGSOFT-FSE 2008), pages 170-180, 2008.

[RV] Runtime Verification workshop series, e.g., http://rv08.in.tum.de/

[San94] Erik Sandewall. Features and Fluents: The Representation of Knowledge about Dynamical Systems. Clarendon Press, Oxford, 1994.

[SDG09] Sebastian Sardina and Giuseppe De Giacomo. Composition of ConGolog programs. In Proc. of the 21st Int. Joint Conf. on Artificial Intelligence (IJCAI 2009), pages 904-910, 2009.

[SDGLL04] Sebastian Sardina, Giuseppe De Giacomo, Yves Lesperance, and Hector J. Levesque. On the semantics of deliberation in IndiGolog - from theory to implementation. Ann. of Mathematics and Artificial Intelligence, 41(2-4):259-299, 2004.

[Sha97] Murray Shanahan. Solving the Frame Problem: A Mathematical Investigation of the Common Law of Inertia. The MIT Press, 1997.

[SNK+08] J.K. Strosnider, P. Nandi, S. Kumarn, S. Ghosh, and A. Arsanjani. Model-driven synthesis of SOA solutions. IBM Systems Journal, 47(3):415-432, 2008.

[SPM09] Shirin Sohrabi, Nataliya Prokoshyna, and Sheila A. McIlraith. Web service composition via the customization of golog programs with user preferences. In Alex T. Borgida, Vinay Chaudhri, Paolo Giorgini, and Eric Yu, editors, Conceptual Modeling: Foundations and Applications - Essays in Honor of John Mylopoulos, volume 5600 of Lecture Notes in Computer Science, pages 319-334. Springer, 2009.

[Su08] Jianwen Su. Semantic Web Services: Composition and Analysis. Bull. of the IEEE Computer Society Technical Committee on Data Engineering, 31(3), 2008.

[Var96] Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In Faron Moller and Graham Birtwistle, editors, Logics for Concurrency: Structure versus Automata, volume 1043 of Lecture Notes in Computer Science, pages 238-266. Springer, 1996.

[vdAAtH+09] W.M.P. van der Aalst, M. Adams, A.H.M. ter Hofstede, M. Pesic, and H. Schonenberg. Flexibility as a Service. In L. Chen, editor, Database Systems for Advanced Applications (DASFAA 2009), volume 5667 of Lecture Notes in Computer Science, pages 320-334. Springer-Verlag, Berlin, 2009.

[vdABEW01] Wil M. P. van der Aalst, Paulo Barthelmess, Clarence A. Ellis, and Jacques Wainer. Proclets: A framework for lightweight interacting workflow processes. Int. J. of Cooperative Information Systems, 10(4):443-481, 2001.

[vdADO+08] W.M.P. van der Aalst, M. Dumas, C. Ouyang, A. Rozinat, and H.M.W. Verbeek. Conformance checking of service behavior. ACM Trans. on Internet Technology, 8(3), 2008.

[vdAP06] W.M.P. van der Aalst and M. Pesic. DecSerFlow: Towards a Truly Declarative Service Flow Language. In M. Bravetti, M. Nunez, and G. Zavattaro, editors, International Conference on Web Services and Formal Methods (WS-FM 2006), volume 4184 of Lecture Notes in Computer Science, pages 1-23. Springer-Verlag, Berlin, 2006.

[vdAPS09] Wil M. P. van der Aalst, Maja Pesic, and Helen Schonenberg. Declarative workflows: Balancing between flexibility and support. Computer Science - R&D, 23(2):99-113, 2009.

[vdARW+07] Wil M. P. van der Aalst, Hajo A. Reijers, A. J. M. M. Weijters, Boudewijn F. van Dongen, Ana Karla Alves de Medeiros, Minseok Song, and H. M. W. Verbeek. Business process mining: An industrial application. Information Systems, 32(5):713-732, 2007.

[vdAvDH+03] W.M.P. van der Aalst, B.F. van Dongen, J. Herbst, L. Maruster, G. Schimm, and A.J.M.M. Weijters. Workflow mining: A survey of issues and approaches. Data and Knowledge Engineering, 47(2):237-267, 2003.

[vdAWM04] W.M.P. van der Aalst, A.J.M.M Weijters, and L. Maruster. Workflow mining: Discovering process models from event logs. IEEE Trans. on Knowledge and Data Engineering, 16(9):1128-1142, 2004.

[vDdMV+99] B.F. van Dongen, A.K.A. de Medeiros, H.M.W. Verbeek, A.J.M.M. Weijters, and W.M.P. van der Aalst. The ProM framework: A new era in process mining tool support export. In S. Donatelli and J. Kleijn, editors, Proc. of the Int. Conf. on Application and Theory of Petri Nets, pages 444-454, 1999.

[Ver09] John Vergo. Artifact-based transformation of IBM Global Financing, a case study. In Proc. of 7th Int. Conference on Business Process Management (BPM 2009), 2009.

[Via09] Victor Vianu. Automatic verification of database-driven systems: a new frontier. In Proc. of the 12th Int. Conf. on Database Theory (ICDT 2009), pages 1-13, 2009.

[Wag03] G. Wagner. The agent-ob ject-relationship metamodel: towards a unified view of state and behavior. Information Systems, 28(5):475–504, 2003

[WB98] P. Wolper and B. Boigelot. Verifying systems with infinite but regular state spaces. In Proc. of the 10th Int. Conf. on Computer Aided Verification (CAV'98), pages 88-97, 1998.

[WJK00] Michael Wooldridge, Nicholas R. Jennings, David Kinny: The Gaia Methodology for Agent-Oriented Analysis and Design. Autonomous Agents and Multi-Agent Systems 3(3): 285-312, 2000

[Woo09] Michael Wooldridge. An Introduction to MuliAgent Systems, second edition, 2009.

[WPD+03] P. Wohed, E. Perjons, M. Dumas, and A. H. M. ter Hofstede. Pattern Based Analysis of EAI Languages - The Case of the Business Modeling Language. In Proceedings of the 5th International Conference on Enterprise Information Systems (ICEIS), pages 174–184, Angers, France, 2003.

[WWG01] G. Wirtz, M. Weske and H. Giese. The OCoN Approach to Workflow Modeling in Object-Oriented Systems. Information Systems Frontiers *3*(3), 357–376, 2001

[WZ99] Frank Wolter and Michael Zakharyaschev. Temporalizing description logic. In D. Gabbay and M. de Rijke, editors, Frontiers of Combining Systems, pages 379-402. Studies Press/Wiley, 1999.

[ZHK96] Gang Zhou, Richard Hull, and Roger King. Generating data integration mediators that use materializations. J. of Intelligent Information Systems, 6:199-221, 1996.