



Grant Agreement N° 215483

Title: Report on industrial dissemination of S-Cube research results

Authors: CNR, FBK, Muenster, POLIMI, SZTAKI, UniDue, UoC, UPM, USI, USTUTT, TUW

Editor: Pierluigi Plebani (POLIMI)

Reviewers: Annapaola Marconi (FBK), Zsolt Nemeth (SZTAKI)

Identifier: CD-IA-2.2.7

Type: Deliverable

Version: 1.0

Date: 29 February 2012

Status: Final

Class: External

Management Summary

The goal of this deliverable is to report on the dissemination during the fourth year of the project to industries which research agenda includes the study on Service Based Applications. Such dissemination has been performed through different channels: individual meetings, workshops, and internships. In addition a specific section on the S-Cube portal has been established to host the publication of the most mature result of the consortium that could be interesting for the companies.

Copyright © 2012 by the S-CUBE consortium – All rights reserved.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n° 215483 (S-Cube).

Members of the S-Cube consortium:

University of Duisburg-Essen (Coordinator)	Germany
Tilburg University	Netherlands
City University London	U.K.
Consiglio Nazionale delle Ricerche	Italy
Center for Scientific and Technological Research	Italy
The French National Institute for Research in Computer Science and Control	France
Lero - The Irish Software Engineering Research Centre	Ireland
Politecnico di Milano	Italy
MTA SZTAKI – Computer and Automation Research Institute	Hungary
Vienna University of Technology	Austria
Université Claude Bernard Lyon	France
University of Crete	Greece
Universidad Politécnica de Madrid	Spain
University of Stuttgart	Germany
University of Hamburg	Germany
VU University Amsterdam	Netherlands

Published S-Cube documents

All public S-Cube deliverables are available from the S-Cube Web Portal at the following URL:

<http://www.s-cube-network.eu/results/deliverables/>

The S-Cube Deliverable Series

Vision and Objectives of S-Cube

The Software Services and Systems Network (S-Cube) will establish a unified, multidisciplinary, vibrant research community which will enable Europe to lead the software-services revolution, helping shape the software-service based Internet which is the backbone of our future interactive society.

By integrating diverse research communities, S-Cube intends to achieve world-wide scientific excellence in a field that is critical for European competitiveness. S-Cube will accomplish its aims by meeting the following objectives:

- Re-aligning, re-shaping and integrating research agendas of key European players from diverse research areas and by synthesizing and integrating diversified knowledge, thereby establishing a long-lasting foundation for steering research and for achieving innovation at the highest level.
- Inaugurating a Europe-wide common program of education and training for researchers and industry thereby creating a common culture that will have a profound impact on the future of the field.
- Establishing a pro-active mobility plan to enable cross-fertilisation and thereby fostering the integration of research communities and the establishment of a common software services research culture.
- Establishing trust relationships with industry via European Technology Platforms (specifically NESSI) to achieve a catalytic effect in shaping European research, strengthening industrial competitiveness and addressing main societal challenges.
- Defining a broader research vision and perspective that will shape the software-service based Internet of the future and will accelerate economic growth and improve the living conditions of European citizens.

S-Cube will produce an integrated research community of international reputation and acclaim that will help define the future shape of the field of software services which is of critical for European competitiveness. S-Cube will provide service engineering methodologies which facilitate the development, deployment and adjustment of sophisticated hybrid service-based systems that cannot be addressed with today's limited software engineering approaches. S-Cube will further introduce an advanced training program for researchers and practitioners. Finally, S-Cube intends to bring strategic added value to European industry by using industry best-practice models and by implementing research results into pilot business cases and prototype systems.

S-Cube materials are available from URL: <http://www.s-cube-network.eu/>

Contents

1	Introduction.....	5
2	S-Cube research results.....	6
2.1	SALMon	6
2.2	JOpera: Process Support for more than Web Services	6
2.3	A service platform for development, deployment and runtime management of Real-Time Online Services	7
2.4	SNAPT.....	8
2.5	Method for observation and simulation of behavior in virtual communities and solutions for management	8
2.6	Runtime Monitoring and Analysis of Influential Factors of Business Process Performance.....	9
2.7	Method for simulation of provision of service orchestrations and resource management policies	9
2.8	A Model Transformation for Increasing Value in Service Networks through Intangible Value Exchanges	10
2.9	Method for the creation of lightweight workflows in social networks.....	10
2.10	Online Testing for Proactive Adaptation	11
2.11	Method for derivation of data-aware upper and lower bounds of QoS attributes for service compositions	11
2.12	SAVVY-WS: A methodology for specifying and verifying web service compositions	12
2.13	Federated SLA aware cloud infrastructures.....	12
2.14	A chemical-based mechanism for workflow instantiation in Service-Based Applications....	13
2.15	A Context-driven Adaptation Process for Service-based Applications.....	14
2.16	Deriving Composite Service Specifications	15
2.17	Method for automatic inference of data attributes of service composition activities and fragments.....	15
2.18	ASTRO – A Framework for Dynamic Composition and Adaptation of SBAs.....	16
3	S-Cube industrial alignment	17
3.1	Meetings.....	17
3.2	Internships.....	18
3.3	Workshops	20
4	Case studies description methodology.....	21
5	Internal initiatives: key recommendations from Industrial Advisory Board	21
6	References.....	24

1 Introduction

The relationship between industries and the S-Cube research community is the main topic of the WP-IA-2.2. After a first phase of the project in which the main objective was to understand the industrial needs, during the last period the workpackage has been more focused on the dissemination of the results to the industries.

Goal of this deliverable is to report the activities on this industrial dissemination of the results achieved by the S-Cube consortium. The main objective of these activities is to promote the results with a potential impact with respect to the interests from the companies. For this reason, particular emphasis has been given to the more mature results that already proposed software implementations or methods and methodologies that have been validated on several scenarios.

To align this dissemination with the industrial needs, the dissemination activities have been coordinated with the members of the S-Cube Industrial Advisory Board that give guidelines to select the results with major impact on the industrial research agenda and that also suggested the way in which these results should be disseminated.

The alignment between the industrial needs and the research community has been supported with several internships: some Ph.D. students of the partners in the S-Cube consortium had the opportunity to directly work in important Research Labs around the world and to bring the competencies inside the NoE.

In addition to the results dissemination, this deliverable also reports on the impact of the methodology to describe the case studies that has been proposed in CD-IA-2.2.2 [2] and improved in CD-IA-2.2.5 [3].

The dissemination activities reported in this deliverable have been performed in coordination with the Spread of Excellence WP. More precisely, the material presented during the industry workshops is also published on the Virtual Campus facilities (WP SoE-1.1), and the dissemination to the industries will be done in collaboration with the ‘Community Outreach’ workpackage (WP SoE-1.2)

The rest of the deliverable is structured as follows. Section 2 summarizes the research results that have been selected as relevant for the industries. Section 3 introduces the dissemination activities performed. Section 4 reports on the results of the dissemination of the case study description methodology. Finally, Section 5 gives details on the IAB meetings.

2 S-Cube research results

The following paragraphs report the 18 research results that constitute the so-called ‘shopping list’ for the industries. This list includes the set of S-Cube research results achieved during the project that companies could have potential interest on.

According to the suggestion given by the Industrial Advisory Board, each result is described not only in terms of objectives and technical details, but also in terms of maturity level (from commercial to paper-based result) and the relationship with the most relevant topics in the industrial research agenda (i.e., cloud computing and future internet applications).

The information reported hereafter constitutes a subset of a complete description of the research results that is available on the S-Cube portal (see <http://s-cube-network.eu/results/s-cube-shopping-list-for-industry-1/results>)

2.1 *SALMon*

Description	SALMon is a monitoring framework to retrieve and assess the Quality of Service (QoS) of services providing different monitoring capabilities. The framework is able to retrieve the QoS on both passive monitoring and testing approaches, check their values, and report the results to the interested parties. SALMon is currently implemented for SOAP-based web services, although it's architecture support other kind of services (e.g. RESTful).
Area	Service Monitoring, Quality of Service
Maturity Level	Prototype for different scenarios. SALMon has been used in several frameworks for the following purposes: NFP-based service selection, self-healing service based systems, quality prediction and monitoring of the cloud.
Relationship with cloud and Future Internet	SALMon is being used, in a current collaboration with the research group of SZTAKI, to test the QoS of the resources and capabilities of the cloud at the infrastructure layer (e.g. network, CPU, etc.). This is achieved by deploying a predefined service with several methods that makes usage of the resources of the cloud. By means of monitoring the service, the QoS of the cloud can be obtained.

2.2 *JOpera: Process Support for more than Web Services*

Description	Eclipse-based Visual Composition Environment featuring a general language for heterogeneous service composition (e.g., RESTful Web Services, WS-* Web Services, Human Services). The language is compiled to Java bytecode for efficient execution within a scalable distributed (i.e., clusters/multicores/cloud) engine architecture. The engine can adapt its configuration autonomously depending on its workload. The engine architecture features an open service invocation mechanism, which can be extended with customized plugins. Likewise, processes can be published as services through different APIs (Hypermedia, WSDL, as well as Java). Human-provided services are composed through a basic, Web-based task worklist handler.
Area	Business Process Management, Model-Driven Service Composition,

Engineering Adaptive Service-Based Systems

Maturity Level Almost Commercial

Relationship with cloud and Future Internet JOpera demonstrates the emerging concept of RESTful business process management. The process execution engine of JOpera has been experimentally deployed in a distributed virtualized environment and tested for elastic scalability properties. Whereas the initial experiment was done in the cloud infrastructure developed within the RESERVOIR EU Project, we plan to further extend the set of cloud environments in which the JOpera engine can be deployed

2.3 *A service platform for development, deployment and runtime management of Real-Time Online Services*

Description The Real-Time Framework (RTF) is a middleware technology for high-level development of scalable Real-Time Online Services along the S-Cube Lifecycle Model. RTF supports three different parallelization and distribution techniques to scale Real-Time Online Services on multiple Cloud servers: zoning, replication and instancing. RTF was used in several industrial applications, e.g.: a multi-server port of the commercial action game Quake 3; a 3D online game Hunter developed by the game company Darkworks; a remote e-learning framework edutain@grid Virtual Classroom developed by the environmental consulting company BMT Cordah Ltd.

RTF is complemented by the Host Management Interface (HMI) that supports the transparent resource management for a running application, in particular the creation, controlling and monitoring of Real-Time Online Services. HMI offers an application providers a management interface for Real-Time Online Services which are implemented on top of RTF. Application providers can deploy their Real-Time Online Services on their own resources or in the Cloud.

Area Software Engineering Life-Cycle, Service adaptation, Grid & Cloud Computing

Maturity Level Prototype (for different scenarios)

Relationship with cloud and Future Internet Future Internet and Internet of Services with their Infrastructure-as-a-Service (IaaS) approach offer new opportunities for the execution of Real-Time Online Services and promise a potentially unlimited scalability by distributing application processing on an arbitrary number of resources given suitable adaptation mechanisms. IaaS allows for adding/removing resources on demand. This opens for Real-Time Online Services an opportunity to serve very high numbers of users and still comply with QoS demands. Despite a variable number of users, IaaS resources can be used efficiently if the application supports adding/removing resources during runtime. Hence, using Cloud Computing for resource provision and the Lifecycle model for implementing adaptable Real-Time Online Services complement each other.

Our middleware platform supports a high-level development approach for Real-Time Online Services and automatically distributes service processing on multiple Cloud resources. We have implemented a resource management system that uses Cloud resources for cost-efficient

and dynamic up- and down-scaling of application sessions using dynamic load balancing based on application-specific monitoring values

2.4 *SNAPT*

Description

A prototype tool for constructing service network models and transforming into initial business process models. Specifically SNAPT allows for: modeling service networks as a set of services and business entities based on a proposed meta-model, contacting quantitative and qualitative analysis of models (VNA analysis), transforming service network models to an initial collaborative business process models, and finally converting service network models to simulation model based on Vensim. SNAPT provides process models inputs for Eclipse BPMN Editor based on the BPMN format and IBM WebSphere Business Modeler V7.0 based on the free-form language supported by the tool. There is no limitation regarding the complexity of service networks. Additionally, SNAPT allows defining and assigning Key Performance Indicators to services in service network models in order to imprint strategic goals and requirements. Currently, SNAPT provides a REST-based KPI repository based on the APQC's Process Classification Framework to allow a SNAPT user to browse and select from 800 KPIs categorized in an industry-based structure, via the KPI editor of his application. Next versions of SNAPT will accommodate the mapping of service network metrics to performance metrics.

Area End-to-end processes in Service Networks, Business Process Management

Maturity Level Prototype (for different scenarios)

Relationship with cloud and Future Internet SNAPT is a unique prototype tool allowing for modeling and analysing existing and emerging service networks enabling business analysts to study and explore the vitality of emerging ones. SNAPT's vision is to bridge the existing gap between business theorists, IT developers and economists by providing a common language and platform for understanding services and service delivery.

2.5 *Method for observation and simulation of behavior in virtual communities and solutions for management*

Description

The goal is to allow service providers to experiment and evaluate their decisions for adaptations previously to deploy them to a real SOA environment hosting virtual communities. Behavior models help to observe and assess the behavior of the members in a virtual community. Observations can be deployed to a testbed and the adaptation strategies based on the models can finally be evaluated for the real world use. As a testbed, Genesis2 testbed generator framework (G2) allowing to design, deploy, control, and monitor Web-services is proposed. The final concept would be a twofold. With a monitoring and controlling infrastructure at the center, on the one side, the real SOA environment can be monitored in a first step. The collected data is processed for two reasons. To identify suspicious behavior of virtual communities and to extract models of behavior samples. The latter are than used for simulation. The simulation environment forms the other side of the concept. Any arising adaptations or changes in configuration can now be tested and observed in the simulation environment first before being

applied to the real system.

Area Adaptation and Monitoring Principles, Techniques and Methodologies for Service-based Applications

Maturity Level Prototype (for different scenarios)

Relationship with cloud and Future Internet In the context of Future Internet, Service Based Applications should be regarded as volatile compositions of a number of possibly independent and autonomous services (i.e. software programs or interfaces to humans) connected through the network and performing a set of functionalities whose integration should fulfill the requirement of the SBA end-user taking into account. In particular with virtual communities such as social networks but also professional expert networks growing constantly, anticipating and identifying unexpected side-effects prior to real deployment is essential for a steady maintainability of the services.

2.6 *Runtime Monitoring and Analysis of Influential Factors of Business Process Performance*

Description The approach deals with process performance management in the context of business processes that are implemented as WS-BPEL service compositions. The process performance is monitored at runtime in terms of a set of key performance indicators (KPIs). If monitoring shows that KPI targets are not met, machine learning techniques are used in order to learn the corresponding influential factors the KPIs depend on. Therefore, decision trees are constructed in an automated fashion based on history monitoring data and are presented to the user. Based on the learned influential factors, adaptation actions can be triggered in order to improve the process performance.

Area Service Composition Layer, Process Monitoring, Business Intelligence, Data Mining

Maturity Level Prototype (for different scenarios)

Relationship with cloud and Future Internet None

2.7 *Method for simulation of provision of service orchestrations and resource management policies*

Description The goal is to allow service providers to experimentally predict future behavior and expected QoS properties of hosted service orchestrations under different input regime assumptions and in combination with different provision resource scaling/elasticity policies. Dynamic models are derived from service orchestrations, and coupled with provision resource scaling/elasticity models within a dashboard-style simulation environment.

Area Service Composition, End-to-End Quality Assurance

Maturity Level Prototype (for single scenario)

Relationship with cloud This simulation-based approach can help service providers in planning

and Future Internet the computation resources for service provisioning, different load-balancing and scaling policies, Service Level Agreements (SLAs) offered to service users, and in contingency planning under different scenarios, thus affecting the overall reliability and flexibility of complex service-based Internet applications. Currently working on modeling elasticity of service provision resources that use cloud infrastructures

2.8 *A Model Transformation for Increasing Value in Service Networks through Intangible Value Exchanges*

Description This work focuses on the identification of some patterns in existing service networks and exploit them to reorganize the network by adding the capability to rapidly react to dynamic environment conditions and to changes in business requirements. The key technology that we are exploiting is the informative value that is given by adding a concept of feedback to every interaction, and by adding a reputation managing system to support the whole infrastructure. The final goal is to increase the total value produced in the systems, especially in terms of knowledge and trust that service network participants may use when interacting with unknown entities, thus at the end they will be able to reorganize their processes in a more “agile” and secure way. We studied first how to apply this methodology in the BPM lifecycle described in the SN4BPM model, and second a case study in which we improve the agility of an existing Car Sharing service network.

A future outcome that will be expected in the near future is to complete the integration into the SN4BPM model and provide a validation of the case study using quantitative techniques for expressing the value in addition to the external service value analysis tools that have been used in the preliminary phase of the research.

Area Service Composition, End-to-End Quality Assurance

Maturity Level Prototype (for single scenario)

Relationship with cloud and Future Internet A scenario of several independent entities providing and using Services is the base of the Future Internet and the Internet of Services. This can be supported by a marketplace of several services that can be used and integrated. The fact that each service may be composed with others and owned by different companies makes it possible to model this scenario using the Agile Service Network formalism that we considered in this work. Similarly to the Agile Service Network we used in our example, we expect that our analysis methodology and the adoption of brokering and matching entities can have a significant impact on the increase in the value also in this context.

No direct relationship with the cloud except for the fact that services are offered by partners in the service network according to the SaaS paradigm

2.9 *Method for the creation of lightweight workflows in social networks*

Description A lightweight coordination and collaboration platform, intertwining contemporary social networking platforms and SOA principles is presented. The approach uses Twitter as a platform for collaborations of human and software services in the context of workflows. Tweets allow

reusing existing infrastructures and tools (e.g., twitter clients on mobile devices) for the communication between services and humans. Simultaneously, social network structures can be exploited originating from Twitter follower networks in order to discover (human and software) resources that are required for the execution of a workflow. One is even able to monitor the execution of workflows with Twitter, simply by following Tweets that represent the execution of a workflow.

Area	Service Composition, Workflows
Maturity Level	Prototype (for single scenario)
Relationship with cloud and Future Internet	As any social platform recently the Twitter network arose to a network with a multitude of users providing and following a plethora of subjects. The demonstrated use for workflows is just on direction in which such well established networks could go. The applications become even more exiting with platforms that are dedicated to the processing of tasks such as crowdsourcing platforms.

2.10 *Online Testing for Proactive Adaptation*

Description	The goal is to enable industry building adaptive systems with proactive capabilities. The contribution is about using online testing to actively collect quality data for predicting failures of services in a service-oriented system. Online testing means testing services in parallel to their normal use and operation. The aim is to achieve accurate failure predictions by increasing the amount of available monitoring data (e.g. QoS data). The idea is to exploit synergies with existing monitoring data. On the one hand, to use monitoring data to build usage models which can be used for test case selection. Thus, services are only tested when there is not sufficient monitoring data. On the other hand, monitoring data is used together with the results of Online testing to predict failures of services in a service-oriented system. As part of the life-cycle of adaptive systems, this prediction can be employed for triggering adaptations and, thus, improve proactive adaptation decisions.
Area	End-to-End Quality Provision and SLA Conformance, Adaption and Monitoring, and Composition and Coordination
Maturity Level	Prototype (for single scenario)
Relationship with cloud and Future Internet	None

2.11 *Method for derivation of data-aware upper and lower bounds of QoS attributes for service compositions*

Description	This method tries to infer upper and lower bounds of QoS metrics as functions of input data, and use them for adaptive QoS-aware matching and predictive monitoring. Automated complexity analysis is applied to Horn clause representations of service choreographies to deduce safe upper and lower approximations of computational complexity as functions of input data. These approximations are combined with
--------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

infrastructure metrics to obtain (predicted) QoS bounds.

Area	Software Engineering Life-Cycle, Service adaptation, Grid & Cloud Computing
Maturity Level	Prototype (for single scenario)
Relationship with cloud and Future Internet	Analyzing the computation cost of service orchestration, in terms of the number of events (e.g. steps, invocations of a component service, loop iterations, etc.) and data sizes (e.g. number of nodes or depth of a XML data structure) can be very useful for proactive adaptation, service binding and match-making under given QoS constraints.

2.12 *SAVVY-WS: A methodology for specifying and verifying web service compositions*

Description	SAVVY-WS aims at providing an integrated approach to support lifelong verification of service compositions. It is based on a specification language, around which design-time and run-time quality assurance activities can be performed. The specification language, based on metric first-order temporal logic, has been designed to express the most common property specification patterns occurring in functional and non-functional requirements of service compositions. The language specifies the required properties in terms of logical formulae, called assumed assertions (AAs). Based on the AAs of all services invoked by the workflow, in turn, the composition may offer a service whose properties can also be specified as formulae, called guaranteed assertions (GAs). At design time a formal verification tool is used to check that a composite service delivers its expected functionality and meets the required quality of service (both specified as GAs), under the assumption that the external services used in the composition fulfill their required interfaces (specified as AAs). Since in service-oriented systems the traditional boundary between design time and run time is blurring, verification, must extend to run time. SAVVY-WS supports continuous verification by transforming specifications into run-time assertions. These assertions are used to check for possible deviations from the correct behavior verified at design time.
--------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Area	Quality Assurance, Service-oriented Software Engineering
Maturity Level	Verification at design time: Mock-up (new prototype will be built in Q1 2012) Verification at run time: Prototype (single scenario)
Relationship with cloud and Future Internet	None

2.13 *Federated SLA aware cloud infrastructures*

Description	A Federated Cloud Management architecture acts as an entry point to cloud federations and incorporates the concepts of metabrokering, cloud brokering and on-demand service deployment. The meta-brokering component provides transparent service execution for the users by allowing the system to interconnect the various cloud broker solutions. Cloud brokers manage the number and the location of the utilized virtual machines for the received service requests. The automatic service
--------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

deployment component is capable of optimizing service delivery by encapsulating services as virtual appliances in order to allow their decomposition and replication among the various IaaS cloud infrastructures. The architecture is aimed at supporting highly dynamic service executions by federating heterogeneous cloud infrastructures in a transparent and autonomous manner.

The work unifies five main aspects and puts them into a single architecture:

- SLA Management: SLA Meta Negotiation, Supporting different SLA concepts, SLA violation propagation
- Meta Brokering: Unified entry point for multiple IaaS systems, Past Performance (metrics) and cost based scheduling
- Cloud brokering: VM lifecycle management, Virtual appliance based queuing
- VM management: Automated virtual appliance distribution and optimization, Virtual machine handling on the different IaaS systems
- Generic Monitoring Registry: Benchmarking different cloud providers.

Area	SLA management, IaaS, cloud federation, virtual appliance
Maturity Level	Mock-up
Relationship with cloud and Future Internet	<p>This work is in alignment with the long term research challenges written in Keith Jeery, Burkhard Neidecker-Lutz (eds.): THE FUTURE OF CLOUD COMPUTING, OPPORTUNITIES FOR EUROPEAN CLOUD COMPUTING BEYOND 2010 in the following aspects</p> <ul style="list-style-type: none"> • global cloud ecosystems (e.g. interoperability) • holistic management and control (e.g. integrate all tiers, address cross-boundary scalability/elasticity/multi-tenancy) • service and application mediation (e.g. customizable products, improved accessibility, composition of higher-value products);

2.14 *A chemical-based mechanism for workflow instantiation in Service-Based Applications*

Description	<p>A nature-inspired mechanism to map activities composing a workflow to actual service implementation is proposed. The mapping is modelled as an evolving and always running chemical-based process that “reacts” to changes in service provisioning. Chemical computing has emerged as an unconventional programming model where a program is described in terms of local interactions (chemical reactions) among molecules floating in a chemical solution. The result of a program is represented by the molecules present in the chemical solution when it reaches an inert state, i.e. when no element in the solution can trigger any reactions. These characteristics allow to equip Service Based Applications with evolutionary and adaptive features that are well-suited to represent and implement autonomous behaviour. In particular, the proposed mechanism allows to model changes in service provisioning regarding both service availability and service non-functional properties (QoS features), as system perturbations that can be processed at any time, and by incrementally affecting the execution state of the system.</p> <p>In order to implement and test the proposed mechanism, a chemical programming environment has been developed named Gamma-Scheme.</p>
--------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

It is a hybrid programming environment for chemical computing merging conventional programming techniques (Scheme) with an unconventional programming model (chemical computing), plus libraries, tools, and advanced programming constructs (Racket Scheme).

Area	Self-Adaptation, Adaptation of Service Compositions, Self-Organising Systems
Maturity Level	Mock-up
Relationship with cloud and Future Internet	<p>In the context of Future Internet, Service Based Applications should be regarded as volatile compositions of a number of possibly independent and autonomous services (i.e. software programs or interfaces to humans) connected through the network and performing a set of functionalities whose integration should fulfil the requirement of the SBA end-user taking into account the .</p> <p>It therefore becomes necessary to organize compositions of services on demand in response to dynamic requirements and changes in service availability and/or service non-functional characteristics.</p> <p>At this end, the chemical-based mechanism provides adaptability to these changes by approaching the composition of services as a decentralized and incremental aggregation mechanism governed by local rules such that environmental changes affecting any part of SBA may be processed at any time.</p> <p>The decentralized chemical-inspired mechanism could be used as a qos-based selection mechanism for Cloud resources in the context of IaaS.</p>

2.15 A Context-driven Adaptation Process for Service-based Applications

Description	<p>The goal is to provide guidelines to support the designers to design and develop adaptable Service Based Applications (SBAs). The approach aims at the identification of the role of the context in the adaptation activities. In particular, a context model is defined considering that the context has various different facets as it includes information ranging from the situation in which users exploit a service-based application to the conditions under which the component services can be exploited. We elaborate on how and when the context should be defined, exploited, and evolved, and on the impact it has on the various activities related to adaptation of service-based applications. Moreover it offers a definition of the associations between the adaptation strategies and the triggers of adaptation mechanisms. We consider different adaptation strategies such as service substitution, re-execution, re-composition, fail, service concretization, re-negotiation, compensation, trigger evolution.</p>
Area	Engineering and design Adaptation and Monitoring
Maturity Level	Paper-based
Relationship with cloud and Future Internet	<p>The approach can be used as a design methodology in a future internet architecture. In fact, service personalization according to context and user preferences is one of the main requirements for future internet.</p> <p>The approach can be tailored to be applied to the self-adaptation of a cloud. Some experiments on this issue are on going</p>

2.16 *Deriving Composite Service Specifications*

Description

While existing service description and composition frameworks attempt to describe service compositions using a variety of composition models ranging from orchestrations to choreographies to Finite State Machines, no attempt has been made to handle the problem of automatically producing specifications for a composite service, based on the specifications of the participating services. Our work aims to provide a thorough and efficient process of automatically deriving composite specifications based on the specifications of the participating services by attempting to deduce the minimum subset of these specifications that needs to be exposed to the service consumer. In order to achieve this, we formulate the derivation for all fundamental control constructs, namely sequential composition, AND-Split/AND-Join, OR-Split/OR-Join and XOR-Split/XOR-Join. By deriving preconditions and postconditions for these constructs, we can derive specifications for any composite service that includes such constructs, by considering the composition schema. The derivation process begins by examining the construct deeper in the schema and gradually moves its way upwards, till the whole composition schema is considered. We also handle loop specification by deriving preconditions and postconditions based on the loop invariant. Finally, the case of handling asynchronous execution is addressed by employing the static single assignment form (SSA) in order to make sure that preconditions are evaluated in the context of the request and not the response.

Area

Formal Models and Languages for QoS-Aware Service Compositions

Maturity Level

Paper-based

Relationship with cloud and Future Internet

The derived specifications offer a complete rich formal service description that gives the opportunity to service providers to accurately describe what they offer and for service consumers to understand exactly what is provided. This is of utmost importance in the Internet of Services since it facilitates service reusability and composability.

2.17 *Method for automatic inference of data attributes of service composition activities and fragments*

Description

The goal is to automatically infer finitely describable, domain-specific characteristics (related to e.g. privacy level or quality) of data used by service compositions activities, based on characteristics of data inputs and composition structure. The approach takes into account service orchestrations that involve complex control structures, such as branches, loops, and sub-workflows. Abstract interpretation-based sharing analysis is applied to Horn clause representations of service choreographies and the results are interpreted in terms of concept lattices (from Formal Concept Analysis) featuring data and activities. The inputs to the analysis and the outputs from it are presented as tables that associate the domain-specific characteristics (attributes) with input data items, activities in the composition, and the intermediate and resulting data items.

Area

Service Composition, Data & Information Related Quality

Maturity Level

Paper-based

Relationship with cloud and Future Internet Service compositions are a powerful and flexible mechanism for implementing more complex, cross-organizational processes based on component services (often provided and managed by a third party), and can, in turn, be used. Analyzing characteristics of data in such complex service systems is deeply related to the issues of data security and privacy, resource usage, fragmentation, and dynamic distributed enactment of complex business processes.

2.18 *ASTRO – A Framework for Dynamic Composition and Adaptation of SBAs*

Description	ASTRO is a platform that includes, and integrates in a coherent framework, advanced methodologies and automated support tools for the design and execution of service-based applications implementing distributed business processes. Covered functionalities include composition, verification, execution, monitoring, and adaptation. ASTRO is continuously extended, both in terms of functionalities and in terms of scenarios, as part of the research and teach transfer activities of FBK. It is also the technology core of SAYservice, a spin-off recently launched by FBK in the areas of Service Oriented Computing and Internet of Services.
Area	Service Composition, Service Adaptation, Service Monitoring, Service Verification
Maturity Level	Prototype for the different modules described in the Technical Information. Some of the modules have been engineered by SAYservice and are now part of SAYservice industrial SOA platform.
Relationship with cloud and Future Internet	Monitoring and adaptation of complex Service-based Applications (SBAs) is one of the most challenging research problems for the Future Internet. SBAs must provide their functionality to potential users with the required/agreed qualities of services, cope with the unreliable network on which they operate, and also deal with the changes in their execution environment, in the partner services with which they interact, and in the users preferences and context. This means that the system must probe the execution to discover all these problems as soon as they materialize, and that it must be able to adapt its behavior to cope with them. ASTRO, supporting the development and operation of dynamic adaptive and context aware service-based systems perfectly meets the needs and challenges of Future Internet.

3 S-Cube industrial alignment

The alignment between S-Cube NoE and the industries requires a two-way communication. On the one side, the S-Cube members need to disseminate their results to the companies that are working in the area of service based applications. On the other side, the companies need to make the network aware about their interests and the most relevant topics.

According to this approach, several actions took place in the last period. Companies have been contacted to present the S-Cube results starting from the shopping list introduced in the previous section. In some cases, see Section 3.1, specific meetings have been organized by the individual partner with some companies to present one result. In some other cases, see Section 3.2, the approach interested the whole project and we organized workshops where different results have been presented.

Moreover, see Section 3.3, some companies open positions for S-Cube Ph.D. students for internships in the company facilities.

3.1 Meetings

Company	NovelTech
Department	N/A
S-Cube result	SNAPT - Service Network Analysis and Prediction Tool
Summary of the meeting	NovelTech is a Greek company specializing in e-government services, requirement analysis and business process management, among other things. NovelTech has expressed interest in SNAPT, the prototype tool for modeling and analyzing service networks. In this context, an initial contact was made for presenting the tool and its capabilities. The presentation took place in early June. During this meeting, NovelTech agreed on using SNAPT in one of its projects concerning the analysis and redesign of procedures undertaken and services offered by the Municipality of Heraklion in Crete, Greece. Specifically, SNAPT will be used to study the network of the City of Heraklion and the analysis methodologies developed within the S-cube project will be employed in order to analyze the performance of processes, including human-centric ones. The general goal is to improve the procedures in terms of response time and service request fulfilments. This collaboration will also contribute to the validation and evolution of UoC's research work. Moreover, NovelTech has developed a Project Management System which is currently deployed in the Municipality of Heraklion and which will be used as the basis for runtime execution data employed in the analysis and redesign of processes.

Company	IBM Germany
Department	N/A
S-Cube result	SNAPT - Service Network Analysis and Prediction Tool
Summary of the meeting	In early September, IBM Germany expressed interest in SNAPT. Christos Nikolaou represented UoC in a meeting held in Germany and gave a detailed presentation of the capabilities of SNAPT. Also, a user manual was prepared and submitted to the IBM faculty. The terms under which UoC and IBM Germany will continue their collaboration will be decided early this year.

3.2 Internships

Collaboration with industries also took place in terms of internships. In line with this approach, IBM Labs hosted 4 S-Cube Ph.D. students in Haifa (Israel) and Almaden (US). Next paragraphs give a short description on these internships.

Ph.D. Student	Waldemar Hummer (TUW)
Hosting institution	IBM Haifa Research Labs, Mount Carmel, 31905 Haifa, Israel
Supervisor	Dr. Orna Raz, Dr. Onn Shehory, Dr. Eitan Farchi (Manager) and Dr. Yaron Wolfsthal (Senior Manager)
Internship period	July 2010 - September 2010
Subject	Methods and Coverage Metrics for Testing of Service-Based Systems
Description of the work	<p>The overall topic of the internship was about methods and coverage metrics for testing of service-based systems. We focused on both the service interface level and the service composition level. Test coverage criteria are important for two reasons: firstly, to obtain a measure of the extent to which a service-based system has been tested, and hence to assess the likeliness that the system will function properly in production use; secondly, coverage criteria are a useful means for generating and executing test cases that aim at testing particular aspects of the service-based system.</p> <p>Coverage on the service interface level means comparing the logged service invocations against the theoretically possible invocations as defined by the schema of the service interface. Full coverage of a service interface means combinatorial testing of all possible inputs, which is often infeasible. Therefore, our approach is based on domain partitioning and allows testers to customize the coverage calculation by defining the subset of the interface that is of interest. For the service composition level, we consider dynamic compositions that define a set of candidate services that are selected and bound at runtime. The aim is to perform integration testing, covering most of the possible service combinations. To avoid combinatorial explosion and to restrict the number of generated test cases, we utilize the novel k-node data flow coverage metric. We present a framework called TeCoS (Test Coverage for Service-based systems) that prototypically implements the discussed contributions. The evaluation comprises various experimentations which have illustrated the feasibility and good scalability of the approach.</p>
Results achieved	<ul style="list-style-type: none"> • Waldemar Hummer, Orna Raz, Onn Shehory, Philipp Leitner, and Schahram Dustdar. Test Coverage of Data-Centric Dynamic Compositions in Service-Based Systems. In 4th International Conference on Software Testing, Verification and Validation (ICST). 2011 • Waldemar Hummer, Orna Raz, and Schahram Dustdar. Towards Efficient Measuring of Web Services API Coverage. In 3rd International Workshop on Principles of Engineering Service-Oriented Systems (PESOS), co-located with ICSE. 2011. • An extended version of the ICST paper is currently under review for a Wiley Journal (Software Testing, Verification and Reliability).

Ph.D. Student	Harald Psailer (TUW)
Hosting institution	IBM - ASR (Almaden Service Research), San Jose, CA, US
Supervisor	Dr. Laura Anderson, Dr. Heiko Ludwig
Internship period	June 2010 - September 2010
Subject	Variation Management
Description of the work	During the stay I had a closer look at existing work in service version and variation management techniques. The result was a differentiation between version and variation management and a more general approach for variation models.
Results achieved	Harald Psailer, Heiko Ludwig, Laura Anderson and Ben Shaw, "Identifying and Managing Variation Scope in Service Management", on submission to ICWS 2012.

Ph.D. Student	Nicolò Maria Calcavecchia (POLIMI)
Hosting institution	IBM Haifa Research Labs, Mount Carmel, 31905 Haifa, Israel
Supervisor	Dr. Ofer Biran
Internship period	September 2010 – December 2010
Subject	Placement optimization of virtual servers in a cloud environment
Description of the work	The research activity carried out during this internship focused on the study and comparison of different optimization techniques for the dynamic virtual server placement in a cloud environment. The problem itself is quite complex, involving multiple constraints, including managerial (security, regulations) resource oriented (capacity, energy), and QoS-oriented (performance, high availability). The objective of this research was to define a reasonable set of criteria by which the quality of placement engines can be compared, and the overall performance of the cloud under them can be assessed. Three different placement techniques were examined: one approach is represented by an heuristic engine that tries to maximize the smallest value of a utility function associated with each virtual server; another approach uses a generic Constraint Satisfaction Programming engine with various optimization goals while a third engine employs a linear optimization model and solves it with a known linear programming tool (ILOG CPLEX).
Results achieved	N. M. Calcavecchia, O. Biran, E. Hadad, Y. Moatti, "VM Placement Strategies for Cloud Scenarios", on submission to IEEE CLOUD 2012

Ph.D. Student	Daniel J. Dubois (POLIMI)
Hosting institution	IBM Haifa Research Labs, Mount Carmel, 31905 Haifa, Israel

Supervisor	Dr. David Breitgand
Internship period	September 2010 – December 2010
Subject	A method and model to support Virtual Machine replication
Description of the work	<p>The purpose of this research work is the following: (i) improve auto-scaling of Virtual Machines by doing rapid clone provisioning using Copy-onWrite; (ii) optimize clone performance on data accesses by doing smart clone pre-fetching; (iii) minimize pre-fetching impact on the virtualization environment. This work is based on the assumption that initial provisioning using CopyOnWrite images is faster than full copy of the VM images and that write accesses to the CopyOnWrite images slow down clone performance. Under these assumptions we studied a way to minimize latency due to the write accesses to the clone, to minimize latency due to read accesses if the Master image is not co-located with the clone, to minimize bandwidth and storage consumption for the cloning operation. The idea behind the solution is to use historical usage information for the data stored in the Master image and then use this information to support the decision on which data need to be prefetched first when a CopyOnWrite image of a clone is being created. This allows the VM clone to be started before all the data are transferred and at the same time reduces the delay for transferring needed data on demand.</p>
Results achieved	<ul style="list-style-type: none"> • A preliminary version of a QEMU (an open source VM emulator and hypervisor) driver that is able to monitor historical data accesses and to apply our prefetching mechanism for our internal experiments. • An internal draft paper explaining the approach has been written, but it has not been submitted yet. David Breitgand, Amir Epstein, Kenneth Nagin, and Daniel J. Dubois, "An Adaptive Online Prefetching Approach to Support Virtual Machines Replication"

3.3 Workshops

On the Feb, 24th, 2012 a workshop in the Thales facilities has been organized with the objective of deepening the following S-Cube research results that Thales considers more relevant with respect to their research agenda:

- SNAPT - Service Network Analysis and Prediction Tool
- Method for simulation of provision of service orchestrations and resource management policies
- Jopera - Eclipse-based Visual Composition Environment featuring a general language for heterogeneous service composition
- A service platform for development, deployment and runtime management of Real-Time Online Services
- Deriving Composite Service Specifications
- Online Testing for Proactive Adaptation
- SALMon -testing&monitoring infrastructure and related technologies
- Federated SLA aware cloud infrastructures

Thales selected these results starting from the ‘shopping list’ introduced in the previous section after several internal meetings that involved managers of the different Thales research divisions. These people also participated to the workshop.

From the S-Cube standpoint, the workshop was a good opportunity to get feedbacks on the research results. In addition, Thales is also available to provide some real business cases to test some of the presented results and to play with some of the prototypes to test the real level of maturity. Finally, Thales recognizes an alignment of these results with respect to the objective of FI-ware infrastructure (<http://www.fi-ware.eu/>) and suggested to apply for the next calls for partnership.

This workshop gave also the possibility to finalize the presentations of these results that now are available on the Virtual Campus facilities.

4 Case studies description methodology

Among the activities performed in the workpackage IA-2.2, the organization of the EU Collaboration Working Group at SSAIE (Service and Software Architectures, Infrastructures and Engineering) gave the opportunity to disseminate the methodology for describing use cases developed in the project. As discussed in the deliverables CD-IA-2.2.2 [2] and CD-IA-2.2.5 [3], S-Cube proposed a methodology for describing case studies [1] that are relevant in the Service and Software research. This methodology has been used to homogeneously describe the case studies that S-Cube proposed during the first year of the project. A further validation of the approach has been conducted at European level. The methodology, indeed, has been presented inside the SSAIE Collaboration Working Group on “Collecting Use Cases” with the aim of having feedbacks from other projects. Based on these feedbacks, a dedicated portal has been developed (see <http://scube-casestudies.ws.dei.polimi.it/>). Goal of this portal is to host a case studies that are proposed by European project.

During the last year, a good result has been achieved, as 3 new case studies are now available on the portal.

- Crisis Management Case Study from PLAY project (<http://www.play-project.eu/>).
- Smart Taxi Case Study from PLAY project (<http://www.play-project.eu/>).
- Bug Resolution Case Study from ALERT project (<http://www.alert-project.eu/>)

Finally, during the SSAIE event that took place on Sept, 28th, 2012 other projects manifested their interest to contribute to the collection of case studies. More precisely:

- Earth Observation (EO) from mOSAIC (<http://www.mosaic-cloud.eu/>)
- Intelligent maintenance from mOSAIC (<http://www.mosaic-cloud.eu/>)
- Social simulations from mOSAIC (<http://www.mosaic-cloud.eu/>)
- Testing Techniques and Tools selection from FITTEST (<http://www.facebook.com/FITTEST>)

A complete description that follows the S-Cube methodology will be provided by the responsible of these projects in the next period.

5 Internal initiatives: key recommendations from Industrial Advisory Board

The Industrial Advisory Board (IAB) for S-Cube has been a valuable source for feedback since the beginning of S-Cube. It is composed of the following members from industry:

- Yaron Wolfsthal, IBM

- Wolfgang Gerteis, SAP
- Pascal Bisson, Thales
- Marquart Franz, Siemens
- Stefano de Panfilis, Engineering Ingegneria Informatica
- Frederic Gittler, HP

The IAB has regularly met with project members, usually in connection with the project's Global Meetings. In the following, a summary report is given on the feedback received during the IAB meetings held during the previous two years of the project. We discuss external initiatives taken by the project and discuss them.

6th IAB Meeting (December 15, 2010 in Ghent, BE)

Participants: Yaron Wolfsthal (IBM), Pascal Bisson (Thales), Marquart Franz (Siemens), Stefano De Panfilis (Engineering), Frederic Gittler (HP); Klaus Pohl, Pierluigi Plebani, Elisabetta Di Nitto, Andreas Metzger (reporter)

Key Recommendations:

- S-Cube is encouraged to ensure that – after providing information to industry – follow-up, joint activities and discussions should be fostered
- Concerning the outreach strategy to industry, S-Cube should identify relevant topics and assets for industry (“shopping list”) and prepare them as a 1-slide presentation to be shared with industry
- Industry workshop @ IBM in Haifa was successful. To foster more intense collaboration, it is suggested that also the industry members present their research areas and how they approach the problems during those workshops; invitation should be extended to all IAB members
- Internships have led to very positive experiences. Specifically, successful results of the internships have been published. Recommendation was to stick to at least 3 months for the duration of internships
- The efforts to establish an IFIP WG on software services have been appreciated by the IAB

7th S-Cube IAB Meeting (June 16, 2011 in Budapest, HU)

Participants: Yaron Wolfsthal (IBM), Pascal Bisson (Thales); Klaus Pohl, Pierluigi Plebani, Andreas Metzger (reporter)

Key Recommendations:

- A 3-way model of internships has been proposed to be followed, including the Technion (IL) as a third entity.
- The IAB would appreciate if S-Cube activities could be continued after the end of the network in one way or another, possibly in the context of the FI PPP
- Industry perceives research outcomes centering on engineering cloud-based systems as very relevant; S-Cube may consider providing a vision towards such cloud-based systems

8th S-Cube IAB Meeting (August 30, 2011 in Essen, DE)

Participants: Yaron Wolfsthal (IBM), Sharon Keidar-Barner (IBM), Eitan Farchi (IBM), Roger Kilian-Kehr (SAP), Clarissa Marquezan, Klaus Pohl (chair), Andreas Metzger (reporter)

Key Recommendations:

- Following up on the recommendation to look at engineering cloud-based systems from the previous IAB, the IAB started discussing possible challenges that need to be faced for collaborative applications on the cloud; especially, this included discussions on how to extend the S-Cube life-cycle models and solutions
- IAB very positively welcomed the ambition to restructure the Virtual Campus and to provide lecture modules on the core S-Cube results online.
- IAB members appreciated the effort of S-Cube to set up an S-Cube Day during ICSE in Zurich

9th S-Cube IAB Meeting (October 26, 2011 in Poznan, PL)

Participants: Pascal Bisson (Thales), Marquart Franz (Siemens), Frederic Gittler (HP), Stefano de Panfilis (Engineering), Yaron Wolfsthal (IBM); Barbara Pernici (co-chair / WP-IA-2.2), Klaus Pohl (chair), Andreas Metzger (reporter)

Key Recommendations:

- The IAB very positively perceived the availability of a “shopping list” of S-Cube results for industry
- The IAB acknowledged that S-Cube was successful in setting up an IFIP working group on “software services”
- The IAB understood that the activities of making the S-Cube use case methodology available to other projects are a very important to ensure sustainability. It was suggested to interlink those activities with a possible FI PPP working group on use case definition.
- Another industry workshop has been scheduled at Thales for end of 2011.
- Internships with IBM have successfully continued.

6 References

- [1] A. Bucchiarone, R. Kazhamiakin, V. Mazza, P. Plebani, Describing Case Studies and Classifying Research Approaches, in S. Dustdar, F. Li, “Service Engineering: European Research Results”, October, 2010, Springer WienNewYork.
- [2] S-Cube Deliverable CD-IA-2.2.2 - Collection of Industrial Best Practices, Scenario and Business Cases.
- [3] S-Cube Deliverable CD-IA-2.2.5 - Final version of the methodology for describing pilot cases.