

SENSORIA: Engineering for Service-Oriented Overlay Computers*

Martin Wirsing¹, Laura Bocchi³, Allan Clark², José Luiz Fiadeiro³,
Stephen Gilmore², Matthias Hölzl¹, Nora Koch^{1, 4}, Philip Mayer¹,
Rosario Pugliese⁵, and Andreas Schroeder¹

¹ Ludwig-Maximilians-Universität München, Germany

² University of Edinburgh, United Kingdom

³ University of Leicester

⁴ F.A.S.T. GmbH, Germany

⁵ Università degli Studi di Firenze, Italy

Abstract. Service-oriented computing is emerging as a new paradigm based on autonomous, platform-independent computational entities (called services) that can be described, published and categorised, and dynamically discovered and assembled, resulting in massively distributed, interoperable, evolvable systems. The IST-FET Integrated Project SENSORIA is developing a novel comprehensive approach to the engineering of service-oriented software systems where foundational theories, techniques and methods are fully integrated in a pragmatic software engineering approach. SENSORIA results comprise modelling and programming primitives for services supported by a mathematical semantics, powerful analysis and verification techniques for behavioural properties and quality of service, and model-based transformation and development techniques. We use one of our case studies, automotive systems, for illustration purposes.

1 Introduction

The last decades have shown tremendous advances in the field of Information Technology (IT). In fact, with online access to vast amounts of information, states offering their services for citizens on the World-Wide Web, and software in control of critical areas such as flight control systems, information systems now lie at the very foundations of our society. But it is the business sector where IT has had the most impact. Driven by the need to stay competitive in an environment of rapid change in global markets and business models as well as local regulations and requirements, organizations now strongly depend on a functional and efficient IT infrastructure which is flexible enough to deal with unexpected changes while still offering stability for business processes and connections to customers, partners, and suppliers.

With the emerging trend of automating complex and distributed business processes as a whole, many organizations face the problem of integrating their existing, already vastly complicated systems to reach yet another layer of sophistication in the interconnection of business value, business processes, and information technology. Many IT

* This work has been partially sponsored by the project SENSORIA, IST-2005-016004.

systems are difficult to adapt and work with; progress in the world of businesses is often impeded by the difficulty of changing the existing IT infrastructure. Service-oriented computing (SOC) is a new approach to software development and integration that addresses these challenges. SOC provides the opportunity for organizations to manage their heterogeneous infrastructures in a coherent way, thus gaining new levels of interoperability and collaboration within and across the boundaries of an organization. In addition, SOC promises new flexibility in linking people, processes, information, and computing platforms.

The IST-FET Integrated Project SENSORIA⁶ is a European Community funded project that develops methodologies and tools for dealing with service-oriented computing. It addresses major problems found in current approaches to SOC and provides mathematically founded and sound methodologies and tools for dealing with the amount of flexibility and interoperability needed in these next-generation infrastructures. SENSORIA aims to support a more systematic and scientifically well-founded approach to *engineering of software systems for service-oriented overlay computers*. At the core of our research is a concept of service that generalises current approaches like Web Services and Grid Computing. SENSORIA provides the scientific foundations for a development process where developers can concentrate on modelling the high-level behaviour of the system and use transformations for deployment and analysis purposes. To this end we investigate programming primitives supported by a mathematical semantics, analysis and verification techniques for qualitative and quantitative properties such as security, performance, quality of service, and behavioural correctness, as well as model-based transformation and development techniques.

In the SeCSE reference model most of the research of SENSORIA is located in the Service Integration Layer, although the research touches both the Semantic Layer and the Infrastructure Layer as well as the crosscutting areas of Management Services, Interoperability and Security. SENSORIA addresses grand challenges from the areas of Service Foundations (dynamic reconfigurable architectures, end-to-end security), Service Composition (composability analysis operators for replaceability, compatibility, conformance; QoS aware service composition; business-driven automated composition), Service Management (self-configuring services), and Service Engineering (design principles for engineering service applications, associating a service design methodology with standard software development and business process modelling techniques).

In the remaining sections, we present the SENSORIA development approach and illustrate its building blocks. Sect. 2 contains an overview of the SENSORIA project and introduces the case study that will be used throughout the paper.

Today, both UML and SCA are commonly used to specify service-oriented systems, either individually or in combination. Therefore we show two alternative approaches to modelling: the SENSORIA UML extension for services and the SCA-inspired SRML in Sect. 3 and Sect. 4, respectively.

When orchestrating service-oriented systems it is often necessary to find trade-offs between conflicting system properties. In Sect. 5 we show how soft constraints can be used to model context-dependent restrictions and preferences in order to perform dynamic service selection.

⁶ <http://www.sensoria-ist.eu/>

To illustrate our approach to qualitative and quantitative analysis we show two of the calculi which are used in the *SENSORIA* development environment but normally hidden from the developers. The *COWS* calculus can be used for performing security modelling and enforcement on service orchestrations; in particular, we add annotations for ensuring confidentiality properties (Sect. 6). The *PEPA* process algebra is used to validate the performance of orchestrations regarding previously defined service-level agreements (Sect. 7).

An expanded version of this paper is available as LMU-IFI-PST technical report 0702 [13].

2 *SENSORIA*

SENSORIA is one of the three Integrated Projects of the Global Computing Initiative of FET-IST, the Future and Emerging Technologies action of the European Commission. The *SENSORIA* consortium consists of 13 universities, two research institutes and four companies (two SMEs) from seven countries⁷.

2.1 *SENSORIA* Approach

The core aim of the *SENSORIA* project is the production of new knowledge for systematic and scientifically well-founded methods of service-oriented software development. *SENSORIA* provides a comprehensive approach to the visual design, formal analysis, and automated deployment of service-oriented applications. The *SENSORIA* techniques enable service engineers to model their applications on a very high level of abstraction using service-oriented extensions of the standard UML or standard business process models, and to transform those models to be able to use formal analysis techniques as well as generate executable code.

The *SENSORIA* techniques and tools are built on mathematical theories and methods. These formal results are complemented by realistic case studies for different important application areas including telecommunications, automotive, e-learning, and e-business. The case studies are defined by the industrial partners to provide continuous practical challenges for the new techniques of services engineering and for demonstrating the research results. This approach is shown in Fig. 1.

An important consideration of the *SENSORIA* project is “scalable analysis for scalable systems.” Our vision is that developers of service-oriented systems can develop at a high level of abstraction with support from tools that analyze the system models, provide feedback about the reliability of the system and possible problems, establish the validity of functional and non-functional requirements, and manage the deployment

⁷ LMU München (coordinator), Germany; TU Denmark at Lyngby, Denmark; FAST GmbH München, S and N AG, Paderborn (both Germany); Budapest University of Technology and Economics, Hungary; Università di Bologna, Università di Firenze, Università di Pisa, Università di Trento, ISTI Pisa, Telecom Italia Lab Torino, School of Management Politecnico di Milano (all Italy); Warsaw University, Poland; ATX Software SA, Universidade de Lisboa (both Portugal); Imperial College London, University College London, University of Edinburgh, University of Leicester (all United Kingdom).

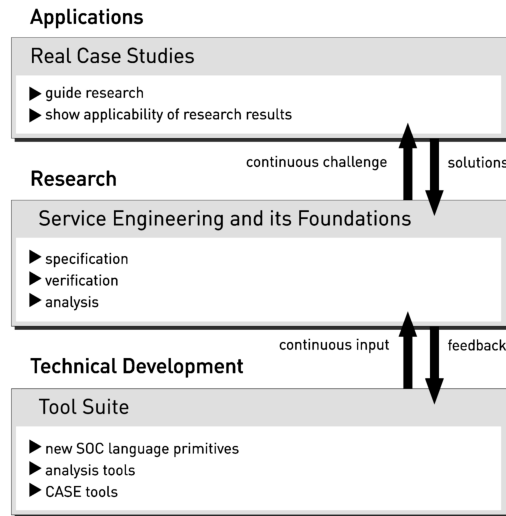


Fig. 1. SENSORIA innovation cycle

to different platforms. We base our tools on formal languages with well-defined properties which makes it possible to establish their correctness. Two aspects are particularly important to make the process practical for commercial software development: (1) analysis results are translated back into the familiar modelling notations so that developers do not need to understand the calculi used to perform the analysis and (2) the analyses are not only usable for “toy examples” but scale to realistic problems. The process may be customized to various application domains and different iterative and agile process models such as the Rational Unified Process (RUP), Model Driven Development (MDD), or Extreme Programming (XP). Fig. 2 displays one step in this development process.

The developer usually works with higher-level input models but also with program code and uses the SENSORIA development environment (downloadable from [12]) to perform qualitative or quantitative analysis and to generate output, e.g., new models or code that can be deployed on various platforms. Our tools within the qualitative and quantitative analyses can, for example, perform checks of functional correctness of services, early performance analysis, prediction of quantitative bottlenecks in collaborating services, and verification of service level agreements.

2.2 Research Themes

The research themes of SENSORIA range across the whole life-cycle of software development from requirements to deployment including reengineering of legacy systems:

- *Modelling Service-Oriented Software.* The definition of adequate linguistic primitives for modelling and programming service-oriented systems enables model-driven development for implementing services on different global computers, and

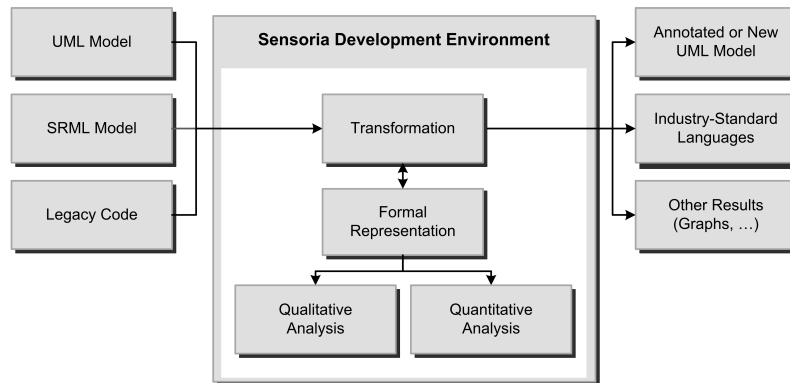


Fig. 2. SENSORIA development process

for transforming legacy systems into services using systematic re-engineering techniques. Modelling front-ends allow designers to use high-level visual formalisms such as the industry standard UML. Automated model transformations allow generation of formal representations from engineering models for further development steps.

- *Formal Analysis of Service-Oriented Software Based on Mathematically Founded Techniques.* Mathematical models, hidden from the developer, enable qualitative and quantitative analysis supporting the service development process and providing the means for reasoning about functional and non-functional properties of services and service aggregates. SENSORIA results include powerful mathematical analysis techniques; in particular program analysis techniques, type systems, logics, and process calculi for investigating the behaviour and the quality of service of properties of global services.
- *Deploying and Runtime Issues of Service-Oriented Software.* The development of sound engineering techniques for global services includes deployment mechanisms with aspects like runtime self-management, service-oriented middlewares, and model-driven deployment as well as reengineering legacy systems into services, thus enabling developers to travel the last mile to the implementation of service-oriented architectures.

2.3 Automotive Case Study

Today, computers embedded in cars can already access communication networks like the Internet and thereby provide a variety of new functionalities for cars and drivers. In the future, instead of having to code these functionalities, services will be able to be discovered at run-time and orchestrated so as to deliver the best available functionality at agreed levels of quality. A set of possible scenarios of the automotive domain are examined to illustrate the scope of the SENSORIA project. In the following we focus on one of them, the *car repair scenario*.

In this scenario, the diagnostic system reports a severe failure in the car engine implying that the car is no longer drivable. The in-vehicle repair support system is able to orchestrate a number of services (garage, back-up car rental, towing truck) that are discovered and bound at that time according to levels of service specified at design time, e.g. balancing cost and delay. The owner of the car deposits a security payment before the discovery of the services is triggered.

3 UML Extension for Service Oriented Architectures

Within the SENSORIA approach, a service engineer can model services using a specialisation of the UML covering static, dynamic, and quality-of-service aspects of services.

For the static aspects of service-oriented software systems, the UML specialisation ranges from rather simple, stereotyped language extensions for introducing services to more complicated structures like dependency relations between services and their contextual relationships to resources and legacy systems. Modeling dynamics of service-oriented software, in particular orchestration and choreography of services, is supported by primitives for interaction and activity modelling that take into account possible failures and quality-of-service aspects (these non-functional extensions are not covered in this paper. Interested readers are referred to [14]).

3.1 Modelling Structures of SOA

The structure of a service-oriented architecture can be visualised by UML deployment and composite structure diagrams. A deployment diagram is used to represent the—usually nested—nodes of the architecture, i.e. hardware devices and software execution environments. Fig. 3 shows a UML deployment diagram of the car and its environment as first approximation to an architecture model. The nodes are connected through communication paths that show the three types of communication that characterise the automotive domain: intra-vehicle communication (non-named connections), inter-vehicle communication, and communication among vehicle and environment such as the communication with the car manufacturer or a remote discovery server. The components that are involved in the execution of service orderings are a service discovery which may be local or external to the car, a reasoner for service selection and a service orchestrator.

In addition to UML deployment diagrams, which give a static view of the architecture, we use UML structure diagrams to represent the evolving connections within the service-oriented architecture of the vehicle and its environment. Three different types of connections are identified: discovery connection, permanent connection (as in modelling of non service-oriented architectures) and temporary connections. In order to be invoked services need to be discovered before the binding to the car's on-board system takes place. Thus the discovery connection is based on the information provided by a discovery service. We distinguish a temporary connection which is for example one from the car on-board system to a known service such as the car manufacturer discovery service, from a permanent connection. Permanent connections wire components as in traditional software. For a graphical representation and more details about these connections the reader is referred to [14].

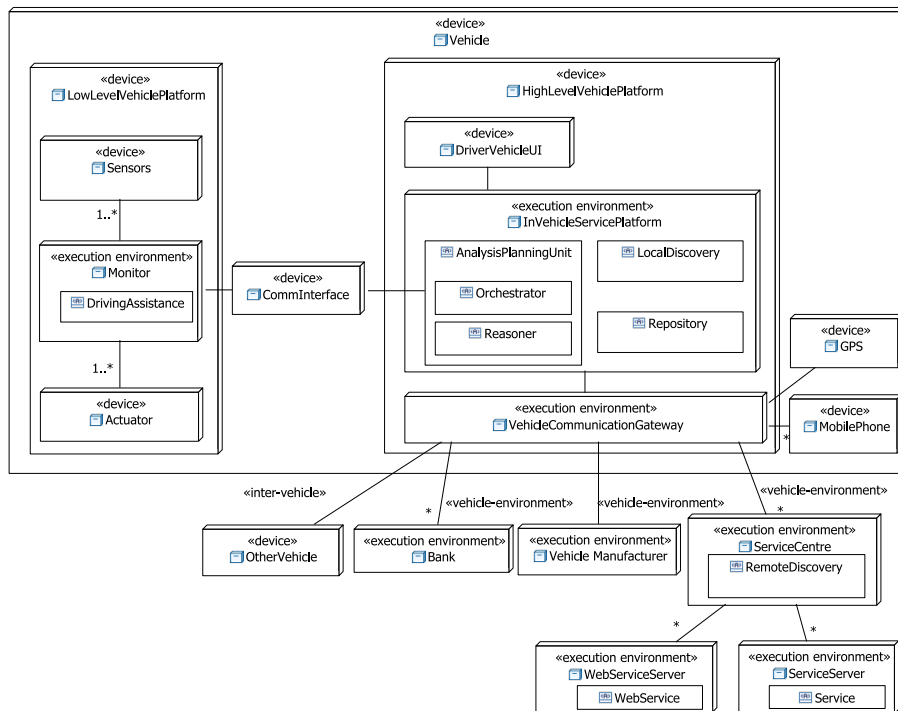


Fig. 3. Simplified architecture of car and car environment

3.2 Modelling Behaviour of SOA

The behaviour of a service-oriented system is mainly described by the service orchestration which defines the system's workflow. Modelling orchestration of services includes specifying interactions among services, modelling transactional business processes using concepts developed for long running transactions, and also specifying non-functional properties of services.

In the modelled business process of the on road car repair scenario the orchestration is triggered by an engine failure or a sensor signal like low oil level. The process starts with a request from the Orchestrator to the Bank to charge the driver's credit card with the security deposit payment, which is modelled by an asynchronous UML action requestCardCharge for charging the credit card the number of which is provided as output parameter of the UML stereotyped call action. In parallel to the interaction with the bank the orchestrator requests the current position of the car from the car's internal GPS service. The current *location* is modelled as input to the requestLocation action and subsequently used by the findLocalServices and findRemoteServices interactions which retrieve a list of services. For the selection of services the Orchestrator synchronises with the Reasoner to obtain the most appropriate services. Service ordering is modelled by the UML actions orderGarage, orderTowTruck and orderRentalCar following a sequential and parallel process, respectively. Fig. 4 shows an UML activity diagram

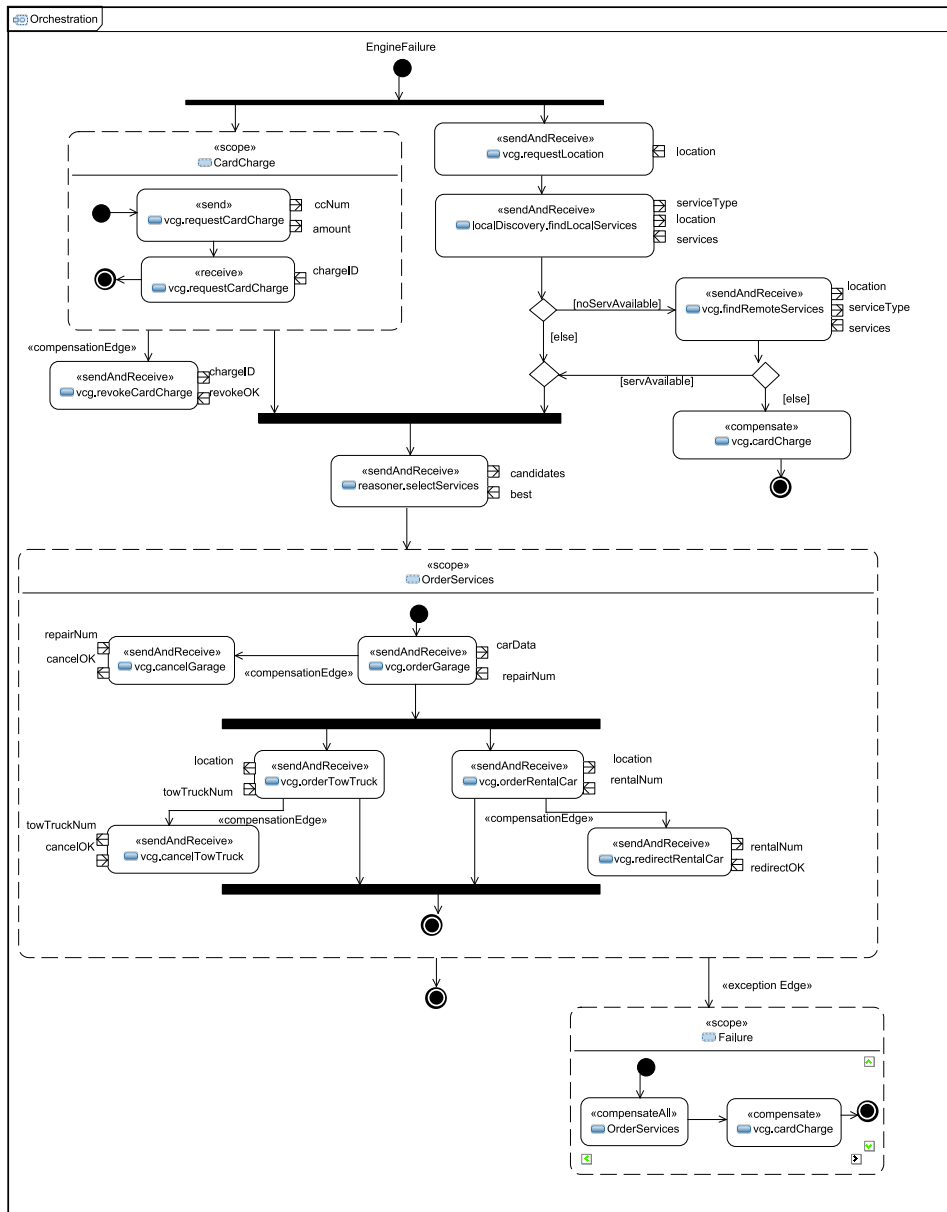


Fig. 4. Orchestration in the on road car repair scenario

of the orchestration of services in the on road car repair scenario. We use stereotyped UML actions indicating the type of interactions (send, receive, sendAndReceive) and model input and output parameters of the interactions with UML pins stereotyped with

outgoing and incoming arrows (abbreviations for *send* and *receive* respectively). These pins are not linked to any edge, but specify variables containing data to be sent or target variables to accept the data received, while constraints prohibit illegal combinations such as e.g. *send* actions with input pins. Interactions match operations of required and provided interfaces of the services. Services are defined as ports of UML components.

The key technique to handle long running transactions is to install compensations, e.g. based on the Saga concepts [6]. Modelling of compensations is not directly supported in UML. We provide an UML extension within SENSORIA for the modelling of these compensations. The extension consists of two modelling primitives – Scope and CompensationEdge – and corresponding stereotypes for UML activity diagrams. A Scope is a structured activity node that groups activity nodes and activity edges. The grouping mechanism provided is convenient for the definition of a compensation or fault handler for a set of activities; at the same time, fault and compensation handlers can also be attached directly to actions. Compensation handlers are defined linking an activity or scope with a compensationEdge stereotyped activity edge to the compensation handler. The handler in turn may again consist of a single action, e.g., see *cancelGarage* in Figure 4. On completion of an activity node with an associated compensation handler the handler is installed and is then executed upon the occurrence of an unhandled exception in the continued execution of the orchestration.

The UML stereotypes defined for orchestration are part of the SENSORIA UML Profile. The use of such an UML extension has direct advantages for the design of model transformations, especially for deployment transformations.

4 The SENSORIA Reference Modelling Language

SRML is a language for modelling composite services understood as services whose business logic involves a number of interactions among more elementary service components as well the invocation of services provided by external parties.

SRML offers modelling primitives inspired by the Service Component Architecture (SCA) [1] but addressing a higher level of abstraction in which one models the business logic of the domain. SRML models both the structure of composite services and their behavioural aspects. As in SCA, interactions are supported on the basis of service interfaces defined in a way that is “independent of the hardware platform, the operating system, the hosting middleware and the programming language used to implement the service” [7].

A SRML module declares one or more components, a number of requires-interfaces that specify services that need to be provided by external parties, and (at most) one provides-interface that describes the properties of the service that is offered by the module. Components have a tight coupling (performed at design time) and offer a distributed orchestration of the parties involved in the service. The coupling between components and external parties is looser and is established at run time through the discovery, selection and binding mechanisms that are supported by the underlying service middleware.

A number of wires establish interaction protocols among the components and between the components and the external parties that instantiate the interfaces. A wire establishes a binding between the interactions that both parties declare to support and

defines the interaction protocol that coordinates those interactions. Figure 5 illustrates the SRML module for the on road car repair scenario.

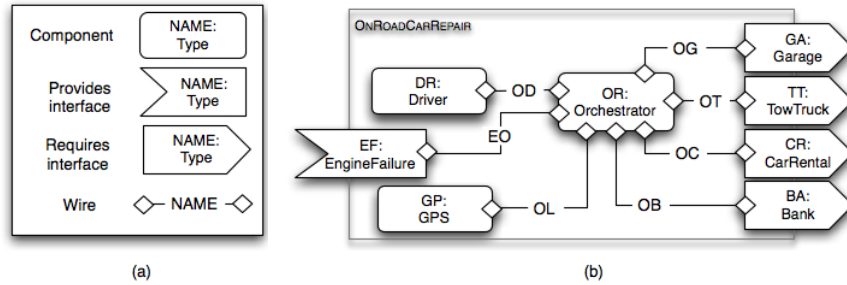


Fig. 5. (a) Entities in a SRML module, (b) SRML module for the on road car repair scenario

We are developing an editor for SRML based on Eclipse Modelling Framework (EMF). The editor relies on a SRML meta-model consisting of a EMF tree; graph transformation techniques can be used to automate the encoding between (parts of) SRML and other languages for which an EMF meta-model exists. Thus, for example, we are providing an encoding from BPEL processes into SRML modules [4]. More specifically, we extract high-level declarative descriptions of BPEL processes that can be used for building more complex modules, possibly including components defined in other (implementation or modelling) languages for which an encoding into SRML exists. We are using the same approach to develop mappings between SRML and other modelling languages/calculi developed in SENSORIA. For instance, we are encoding SRML into COWS in order to provide SRML with an operational semantics for the process of service discovery.

4.1 Specifying SRML entities

The specification of each SRML entity (components, external interfaces and wires) addresses properties of the interactions supported by that entity.

Figure 6 presents the interactions specified in the business role fulfilled by the orchestrator. SRML supports conversational interaction types – send-and-receive (**s&r**) and receive-and-send (**r&s**) – involving initiation ($interaction \begin{smallmatrix} \curvearrowright \\ \text{event} \end{smallmatrix}$ denotes the event of initiating *interaction*), reply ($interaction \begin{smallmatrix} \boxtimes \\ \text{event} \end{smallmatrix}$ denotes the reply-event of *interaction*) and other events for handling specific aspects of conversations such as committing, cancelling and revoking a deal (see [7] for an exhaustive list of types of interaction and interaction-event). Interactions can have $\begin{smallmatrix} \curvearrowright \\ \text{parameter} \end{smallmatrix}$ -parameters for transmitting data when the interaction is initiated and $\begin{smallmatrix} \boxtimes \\ \text{parameter} \end{smallmatrix}$ -parameters for carrying a reply.

The different types of entities involved in a module are specified in SRML using three different but related languages. All the languages provide a specification of the supported interactions. The languages for defining business roles, business protocols and interaction protocols differ in the way they define the behavioural interface.

```

INTERACTIONS
rcv engineFailure
s&r AskUsrDetails
    ☒ cust:customerID
    ccID:bankAccountID
    dest:location
    workrelated:boolean
    Δt:Time
s&r currentLocation
    ☒ l:Location

s&r requestCharge
    ☒ cust:customerID
    ccID:bankAccountID
    amount:moneyVal
    ☒ certif:moneyVal
s&r orderGarage
...
s&r rentACar
...
s&r orderTowTruck
...

```

Fig. 6. Specification of the interactions for the business role *Orchestrator*

A component is a computational unit represented in a module by what we call a business role. A business role declares the interactions in which the component can become involved, and the execution pattern that orchestrates those interactions (i.e., the orchestration). The orchestration is specified through a set of variables that provide an abstract view of the state of the component and by a set of transition rules.

The transition rules can be derived from UML sequence or activity diagrams and refined with additional information. Each transition has a *trigger*, typically the occurrence of an interaction event, a *guard*, identifying the states in which it can take place, the *effects* that change the local state (s' denotes the value of the local variable s after the transition) and the *sends* which specifies the interaction events that are sent. The following is a transition of the business role *Orchestrator*.

```

transition startProcess
  triggeredBy engineFailure!
  guardedBy s=Begin
  effects s' = FailureDetected
  sends currentLocation!
  ^ askUsrDetails!

```

The external interfaces are defined in a module through what we call business protocols. Business protocols specify the interactions similarly to business roles. The behaviour of business protocols abstracts from details of the local state of the co-party and specifies the protocol that the co-party adheres to as a set of properties concerning the causality of the interactions. The following is the behaviour for the business protocol *Customer*.

```

initiallyEnabled engineFailure!
engineFailure! ensures confirmation!

```

4.2 Service Level Agreement

SRML offers primitives for modelling “dynamic” aspects concerned with configuration, session management and Service Level Agreement (SLA). In particular, SRML supports the definition of attributes and constraints for SLA using the algebraic techniques developed in [3,2] for constraint satisfaction and optimization. Section 5 presents this calculus in more detail together with an example. The constraints presented therein can be modeled in SRML using a constraint system as described in the following section.

5 Soft Constraints for Selecting the Best Service

The reasoning component of a service-oriented system (cf. Fig. 3) has to achieve a compromise between different goals of the system. Soft constraints are a promising way to specify and implement such reasoning mechanisms.

Soft constraints are an extension of classical constraints to deal with non-functional requirements, over-constrained problems and preferences. Instead of determining just a subset of admissible domain elements, a soft constraint assigns a grade to each element of the application domain. Bistarelli, Montanari and Rossi [3,2] have developed a very elegant theory of soft constraints where many different kinds of soft constraints can be represented and combined in a uniform way over so-called constraint semirings (*c-semirings*).

In SENSORIA we are developing a language which extends the *c-semiring* approach with possibilities to specify preferences between constraints and to vary the constraints according to a *context*. This simplifies the description of behaviours of systems in a dynamically changing environment and with complex trade-offs between conflicting properties.

A context is a boolean expression which can guard a constraint. For example, the distance to the destination might determine whether the quick availability of a rental car is important or not. In this case, “*distance < 20km*” is a context that can restrict the applicability of a constraint. Variables appearing in contexts are called *context variables*; variables appearing only in constraints but not in contexts are called *controlled variables*. In the car repair scenario the context variables will specify, among others, the *distance* to our destination or whether the journey is *work-related*. The controlled variables represent properties of offers, e.g., the cost or quality of an offer.

A soft constraint is a mapping from (domains of) controlled variables into a *c-semiring*. In the car repair scenario we use soft constraints to specify the preferences of the users. For example, the user may prefer garages that can repair the car as quickly as possible with durations of more than two days being not acceptable:

$$fastRepair : [garage-duration \mid n \mapsto \lfloor 48/n \rfloor]$$

We also may want the repair to be done cheaply, but only if the trip is not work related. Repairs costing more than 1000 Euros are still acceptable, but only barely:

$$cheapRepair : \mathbf{in\ context} \neg work-related? \\ \mathbf{assert} [garage-cost \mid n \mapsto \lceil 1000/n \rceil] \mathbf{end}$$

Users might not consider all constraints to be equally important. For example a user who prefers fast but expensive repairs can specify this as preference between constraints: $fastRepair > cheapRepair$. In this case an offer that results in a faster repair will always be preferred; the price will only be used to break ties between equally fast repairs.

From a set of constraints and preferences the reasoner can compute either the best solutions, or a set of all solutions that are better than a certain threshold. Two techniques that are used for solving soft constraint systems are branch and bound search [16] and dynamic programming [2].

The constraints presented in this section can be modeled in SRML using the following syntax where S is a c-semiring, D is a finite set (domain of possible elements taken by the variables), and V is a totally ordered set (of variables):

CONSTRAINT SYSTEM

```
S is <[0..1000],max,min,0,1000>
D is { n∈ℕ:1≤ n ≤1000 }
V is { DR.askUsrDetails⊗.workRelated, GA.cost, GR.duration,...}
```

The variable $DR.askUsrDetails⊗.workRelated$ is a Boolean context variable that is true if the trip is work related. The variables $GA.cost$, $GR.duration$ represent the cost and availability of the garage service, respectively. We show the SRML specification of the *fastRepair* and *cheapRepair* constraints. Notice that we represent the context as a variable of the constraint.

CONSTRAINTS

```
fastRepair is <{GA.duration},def1> s.t. def1(n)=floor(48/n)
cheapRepair is <{GA.cost},def2> s.t. def2(n)=
  if DR.askUserDetails⊗.workrelated
  then 1000 else ceiling(1000/n)
```

6 A Process Calculus for Service-Oriented Systems

COWS (*Calculus for Orchestration of Web Services*, [11]) is a recently designed process calculus for specifying, combining and analyzing service-oriented applications, while modelling their dynamic behaviour. We present (an excerpt of) COWS main features and syntax while modelling some simplified components of the on road car repair scenario. The type system of COWS [10] enables us to verify confidentiality properties, e.g., that critical data such as credit card information is shared only with authorized partners. The complete specification, including compensation activities, can be found in [13].

6.1 Service Orchestration with COWS

To start with, we present the COWS term representing the ‘orchestration’ of all services within the scenario of Sect. 3:

$$[p_{car}](Orchestrator \mid LocalDiscovery \mid Reasoner \mid SensorsMonitor) \mid Bank \mid OnRoadRepairServices$$

The services above are composed by using the *parallel composition* operator \mid that allows the different components to be concurrently executed and to interact with each other. The *delimitation* operator $[_]$ is used here to declare that p_{car} is a (partner) name known to all services of the in-vehicle platform, i.e. *Orchestrator*, *LocalDiscovery*, *Reasoner* and *SensorsMonitor*, and only to them.

Orchestrator, the most important component of the in-vehicle platform, is

$$[x_1, \dots, x_n] (p_{car} \bullet o_{engfail}? \langle x_1, \dots, x_n \rangle . s_{engfail} \\ + p_{car} \bullet o_{lowoil}? \langle x_1, \dots, x_n \rangle . s_{lowoil})$$

This term uses the *choice* operator $_+_$ to pick one of those alternative ‘recovery’ behaviours whose execution can start immediately. For simplicity, only ‘engine failure’ and ‘low oil’ situations are taken into account.

The *receive-guarded prefix* operator $p_{car} \bullet o_i? \langle x_1, \dots, x_n \rangle .$ expresses that each recovery behaviour starts with a *receive* activity of the form $p_{car} \bullet o_i? \langle x_1, \dots, x_n \rangle$ corresponding to reception of a request emitted, when a failure arises, by *SensorsMonitor* (a term representing the behaviour of the ‘low level vehicle platform’ of Fig. 3). *Receives*, together with *invokes*, written as $p \bullet o! \langle e_1, \dots, e_m \rangle$, are the basic communication activities provided by COWS. Besides input parameters and sent values, they indicate an *endpoint*, i.e. a pair composed of a partner name p and an operation name o , through which communication should occur. $p \bullet o$ can be interpreted as a specific implementation of operation o provided by the service identified by the logic name p . An inter-service communication takes place when the arguments of a receive and of a concurrent invoke along the same endpoint do match, and causes replacement of the variables arguments of the receive with the corresponding values arguments of the invoke (within the scope of variables declarations). For example, variables x_1, \dots, x_n , declared local to *Orchestrator* by means of the delimitation operator, are initialized by the receive leading the recovery activity with data provided by *SensorsMonitor*.

The recovery behaviour $s_{engfail}$ executed when an engine failure occurs is

$$[p_e, o_e, x_{loc}, x_{list}] \\ ((requestCardCharge \mid requestLocation. findServices) \\ \mid p_e \bullet o_e? \langle \rangle . p_e \bullet o_e? \langle \rangle . selectServices. orderGarage. \\ (orderTowTruck \mid orderRentalCar))$$

$p_e \bullet o_e$ is a scoped endpoint along which successful termination signals (i.e. communications that carry no data) are exchanged to orchestrate execution of the different components. Variables x_{loc} and x_{list} are used to store the value of the current car’s GSP position and the list of closer on road services discovered. To present the specification of $s_{engfail}$ in terms of the UML actions of Fig. 4, we have used an auxiliary ‘sequence’ notation as e.g. in $requestLocation. findServices$. This notation indicates that execution of $requestLocation$ terminates before execution of $findServices$ starts. Indeed, $requestLocation. findServices$ actually stands for the COWS term

$$p_{car} \bullet o_{reqLoc}! \langle \rangle \mid p_{car} \bullet o_{respLoc}? \langle x_{loc} \rangle . \\ (p_{car} \bullet o_{findServ}! \langle x_{loc} \rangle \\ \mid p_{car} \bullet o_{found}? \langle x_{list} \rangle . p_e \bullet o_e! \langle \rangle + p_{car} \bullet o_{notFound}? \langle \rangle)$$

where $requestLocation$ and $findServices$ are

$$requestLocation \triangleq p_{car} \bullet o_{reqLoc}! \langle \rangle \mid p_{car} \bullet o_{respLoc}? \langle x_{loc} \rangle \\ findServices \triangleq p_{car} \bullet o_{findServ}! \langle x_{loc} \rangle \\ \mid p_{car} \bullet o_{found}? \langle x_{list} \rangle . p_e \bullet o_e! \langle \rangle + p_{car} \bullet o_{notFound}? \langle \rangle$$

Bank, the last service we show, can serve multiple requests simultaneously. This behaviour is modelled by exploiting the *replication* operator $*_*$ to spawn in parallel as many copies of its argument term as necessary. The definition of *Bank* is

$$\begin{aligned}
& * [x_{cust}, x_{cc}, x_{amount}, o_{checkOK}, o_{checkFail}] \\
& p_{bank} \bullet o_{charge} ? \langle x_{cust}, x_{cc}, x_{amount} \rangle \cdot \\
& (\langle \text{perform some checks and reply on } o_{checkOK} \text{ or } o_{checkFail} \rangle \\
& \quad | p_{bank} \bullet o_{checkOK} ? \langle \rangle \cdot x_{cust} \bullet o_{chargeOK} ! \langle \rangle \\
& \quad + p_{bank} \bullet o_{checkFail} ? \langle \rangle \cdot x_{cust} \bullet o_{chargeFail} ! \langle \rangle)
\end{aligned}$$

Once prompted by a request, differently from *Orchestrator*, *Bank* creates one specific instance to serve that request and is immediately ready to concurrently serve other requests. Notably, each instance exploits communication on ‘internal’ operations $o_{checkOK}$ and $o_{checkFail}$ to model a conditional choice.

6.2 Using Types for Verifying Service Properties

One advantage of using COWS as modelling language is that it already provides some tools for analysing the models and verifying the properties they enjoy. For example, the type system introduced in [10] for checking data security properties on COWS terms is a practical and scalable way to provide evidence that a large number of applications enjoy some given properties: from the *type soundness* of the language as a whole, it follows that all *well-typed* applications do comply with the properties stated by their types. The types permit to express policies constraining data exchanges in terms of *regions*, i.e. sets of partner names attachable to each single datum. Service programmers can thus settle the partners usable to exchange any given datum (and, then, the services that can share it), thus avoiding the datum be accessed (by unwanted services) through unauthorized partners. The language operational semantics uses these annotations to guarantee that computations proceed according to them.

To provide a flavour of the properties that can be expressed and enforced by using the type system, we illustrate some properties relevant for the scenario modelled before. Firstly, a driver in trouble must be assured that information about his credit card and GSP position cannot become available to unauthorized users. To this aim, the credit card identifier $ccId$, communicated by activity $requestCardCharge$ to service *Bank*, can be annotated with the policy $\{p_{bank}\}$, that allows *Bank* to receive the datum but prevents it from transmitting the datum to other services. Similarly, the car’s GSP position stored in x_{loc} , used by services $orderGarage$, $orderTowTruck$ and $orderRentalCar$, can be annotated with the regions $\{x_{garage}\}$, $\{x_{towTruck}\}$ and $\{x_{rentalCar}\}$, respectively, to specify different policies for the same datum, according to the invoked services. Notably, these policies are not statically fixed, but *depend* on the partner variables x_{garage} , $x_{towTruck}$ and $x_{rentalCar}$, and, thus, will be determined by the values that these variables assume as computation proceeds. As a final example property, we mention that, by using appropriate regions only including the customer partner name, one can guarantee that critical data of on road services, such as cost and quality of the service supplied, are not disclosed to competitor services.

7 Quantitative Analysis of Service Level Agreements

In the car repair scenario the quantitative issue of concern relates to how long it will take from the point of engine failure until both a tow truck and a garage have been ordered, and the tow truck is on its way to help the stranded driver. If the duration of each of the service activities which need to take place along the way (*requestLocation*, *findServices*, *orderGarage*, ...) was known exactly then this calculation would simply involve adding up these times to give the total duration. However, as is usual in service-oriented systems, none of these durations will be known exactly here and the calculation needs to be based on the expected average duration of each of the atomic service events. In this setting, where only the average duration is known, and not the variance or higher moments, then the correct distribution to model with is the exponential distribution. Thus, this aspect of the problem naturally lends itself to Markovian representation and analysis, and that is the approach which we will take here.

For the quantitative analysis of such systems we use Performance Evaluation Process Algebra (PEPA) [9] which extends classical process algebras by associating a duration with each activity. Thus where classical process algebras such as CCS and CSP deal with instantaneous actions which abstract away from time, PEPA has instead continuous-time activities whose durations are quantified by exponentially-distributed random variables. Thus PEPA is a *stochastic* process algebra which describes the evolution of a process in continuous time. The operational semantics of the language gives rise to a continuous-time finite-state stochastic process called a Continuous-time Markov Chain (CTMC) which can be used to find the steady-state probability distribution over a model. From this it is straightforward to compute conventional performance measures such as utilisation or throughput. Here we are instead performing transient analysis of the CTMC where one considers the probability distribution at a chosen instant of time. It is possible to use these results to perform more complex quantitative analysis such as computing response time measures and first passage time quantiles as used in service-level agreements. It is also possible to perform scalability analysis by using an altogether different representation based on ordinary differential equations. For example, we have investigated with PEPA how models of Web service execution scale with increasing client population sizes (see [15]).

The PEPA process algebra is a compact formal language with a small number of combinators. Components perform activities. Activities have a type and rate specified using *prefix* (\cdot) so that (α, r) denotes performing activity α at rate r and (α, \top) is a partner for this where the other partner in the cooperation determines the rate. Alternative behaviours can be composed in a *choice* ($+$). Parallel composition of components uses CSP-style synchronisation over a set of activity types (\bowtie). Private behaviour can be hidden ($/$).

To perform quantitative analysis we assign exponentially distributed rates to the activities of our model. Fig. 7 explains the meaning of these. In our modelling we were unwilling to assume that we even knew precisely the average duration of the atomic service events and required only that these values lie in a range between some minimum and some maximum average value. The less certainty we have about the actual value of the rate, the wider this range must be.

Activity	Rate	Value		Meaning
		min	max	
EngineFailure	r_1	1.0	1.0	arbitrary value — measurement only begins at the end of this event
RequestLocation, FindServices	r_2	0.9	1.1	location information can be transmitted in one minute, with little variance, service discovery is similar
CurrentLocation	r_3	0.25	1.25	the driver processes location information and decides to act on this, with high variance
SelectServices	r_4	1.9	2.1	the reasoner takes around thirty seconds to make a decision, with little variance
OrderGarage, OrderTowTruck	r_5	0.25	1.25	the on-road repair service takes slightly less than one minute to process orders, with high variance

Fig. 7. Table of activities and minimum and maximum values of the rates from the model. All times are expressed in minutes. Thus a rate of 1.0 means that something happens once a minute (on average). A rate of 2.0 means that the associated activity happens twice a minute on average, or that its mean or expected duration is thirty seconds, which is an equivalent statement.

We analysed the model first using the PEPA Workbench [8] which confirmed that the model had no deadlocks and no transient states, and that all local states of all components were reachable and that all activities were live.

We used the ipc/Hydra tool chain [5] to assess the service against the following compound service level agreement (SLA-1) on the orchestration overall.

At least 30% of engine failures lead to the tow truck being ordered within fifteen minutes and at least 60% of engine failures lead to the tow truck being ordered within thirty minutes.

At first sight it might seem that SLA-1 is so generous with the time bounds and confidence bounds that it will easily be satisfied and that a more demanding SLA-2 could be posed instead: “30% of failures dealt with in ten minutes, 90% in thirty.” However, there is a possibility that each atomic service event in the orchestration will take longer than average and that each of these is operating at its minimum rate, leading to a much longer overall time than might be expected.

We assess SLA-1 using the passage-time quantile computation capabilities provided by ipc/Hydra. We vary rates r_2 to r_5 across five possible values leading to $5 \times 5 \times 5 \times 5 = 625$ experiments to be performed. The graphs of computed probability against experiment number for time bounds of fifteen minutes and thirty minutes for all 625 experiments are shown in Fig. 8. Using both of these graphs we determine that SLA-1 is satisfied across the values of the rates of the model but that the time bounds and confidence bounds are tight, and the more demanding SLA-2 would *not* be satisfied by some of the combinations of rate values used in the model.

8 Concluding Remarks

Service-oriented computing and service-oriented architecture are having a huge impact on IT-based business organizations across the world. However, there are still many open




Fig. 8. Graph of probability of completing the passage from engine failure to completion of order of tow truck within fifteen and thirty minutes plotted against experiment number over all 625 experiments.

issues regarding the development, analysis, and deployment of such software, some of which touch the very foundations of service-oriented computing.

As a remedy, the EU project SENSORIA is developing a novel comprehensive approach to the visual design, formal analysis, and automated deployment of service-oriented software systems where foundational theories, techniques and methods are fully integrated in a pragmatic software engineering approach.

Compared to other research projects in the field of service oriented computing, SENSORIA focuses heavily on scalable quantitative and qualitative analysis techniques based on formal foundations which can be embedded into a practical software development process. SODIUM⁸ in contrast focuses on service modeling semantic support for service discovery, but leaving aside dynamic reconfiguration, model transformation, service deployment, service extraction, and analysis techniques. PLASTIC⁹ on the other hand is more focused towards development, deployment and management of service oriented adaptive applications. Resource management is a major aspect in the PLASTIC approach. The SeCSE project¹⁰ is geared towards development and runtime support for service-oriented computing, again considering formal analysis techniques of services as a secondary field of activity. SENSORIA's main research focus and asset in the service oriented community is therefore the combination of a pragmatic approach with formal theories for service oriented computing, supporting both development and analysis of service oriented systems.

We have illustrated service-modelling in high level languages like SRML or the SENSORIA extension of UML, service selection with soft constraints, as well as analysis of qualitative and quantitative service properties with process calculi like COWS and PEPA. Further research of the SENSORIA project addresses topics like resource consumption of services, security, reconfiguration, deployment, and re-engineering of legacy systems into services. To facilitate practical application of these results we are distributing the SENSORIA development environment under an open-source license.

⁸ <http://www.atc.gr/sodium/>

⁹ <http://www.ist-plastic.org/>

¹⁰ <http://secse.eng.it/>

To learn more, please visit our website <http://www.sensoria-ist.eu/>.

References

1. Michael Beisiegel, Henning Blohm, Dave Booz, Jean-Jacques Dubray, Adrian Colyer, Mike Edwards, Don Ferguson, Bill Flood, Mike Greenberg, Dan Kearns, Jim Marino, Jeff Mischkinsky, Martin Nally, Greg Pavlik, Mike Rowley, Ken Tam, and Carl Trieloff. Building Systems using a Service Oriented Architecture. Whitepaper, SCA Consortium, 2005. http://www.oracle.com/technology/tech/webservices/standards/sca/pdf/SCA%_White_Paper1_09.pdf.
2. Stefano Bistarelli. *Semirings for Soft Constraint Solving and Programming*. LNCS 2962. Springer, Berlin, 2004.
3. Stefano Bistarelli, Ugo Montanari, and Francesca Rossi. Semiring-based constraint satisfaction and optimization. *J. ACM*, 44(2):201–236, 1997.
4. Laura Bocchi, Yi Hong, Antónia Lopes, and José Luiz Fiadeiro. From BPEL to SRML: A Formal Transformational Approach. In *Proc. of 4th International Workshop on Web Services and Formal Methods (WSFM'07)*, LNCS. Springer, 2007.
5. Jeremy T. Bradley and William J. Knottenbelt. The ipc/HYDRA tool chain for the analysis of PEPA models. In *Proc. 1st Int. Conf. on the Quantitative Evaluation of Systems (QEST 2004)*, pages 334–335, Enschede, Netherlands, September 2004.
6. Roberto Bruni, Gianluigi Ferrari, Hernán Melgratti, Ugo Montanari, Daniele Strollo, and Emilio Tuosto. From theory to practice in transactional composition of web services. In M. Bravetti, L. Kloul, and G. Zavattaro, editors, *Formal Techniques for Computer Systems and Business Processes: Proc. of WS-FM 2005, 2nd Internat. Workshop on Web Services and Formal Methods*, volume 3670 of LNCS, pages 272–286. Springer Verlag, 2005.
7. José Luiz Fiadeiro, Antónia Lopes, and Laura Bocchi. A formal approach to service component architecture. In *Web Services and Formal Methods*, volume 4184 of LNCS, pages 193–213. Springer, 2006.
8. Stephen Gilmore and Jane Hillston. The PEPA Workbench: A Tool to Support a Process Algebra-based Approach to Performance Modelling. In *Proc. of the Seventh Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, volume 794 of LNCS, pages 353–368, Vienna, May 1994. Springer-Verlag.
9. Jane Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
10. A. Lapadula, R. Pugliese, and F. Tiezzi. Regulating data exchange in service oriented applications. In *Proc. of IPM International Symposium on Fundamentals of Software Engineering (FSEN'07)*, volume 4767 of *Lecture Notes in Computer Science*, pages 223–239. Springer, 2007.
11. Alessandro Lapadula, Rosario Pugliese, and Francesco Tiezzi. A calculus for orchestration of web services. In *ESOP*, volume 4421 of LNCS, pages 33–47. Springer, 2007.
12. SENSORIA project. Web site for the SENSORIA development environment. <http://svn.pst.ifi.lmu.de/trac/sct>.
13. Martin Wirsing, Laura Bocchi, Allan Clark, José Luiz Fiadeiro, Stephen Gilmore, Matthias Hözl, Nora Koch, Philip Mayer, Rosario Pugliese, and Andreas Schroeder. SENSORIA: Engineering for Service-Oriented Overlay Computers. Technical Report 0702, Ludwig-Maximilians-Universität München, Institut für Informatik, PST, 2007.
14. Martin Wirsing, Allan Clark, Stephen Gilmore, Matthias Hözl, Alexander Knapp, Nora Koch, and Andreas Schroeder. Semantic-Based Development of Service-Oriented Systems.

- In E. Najm et al., editor, *Proc. 26th IFIP WG 6.1 Internat. Conf. on Formal Methods for Networked and Distributed Systems (FORTE'06), Paris, France*, volume 4229 of *LNCS*, pages 24–45. Springer, 2006.
15. Martin Wirsing, Rocco De Nicola, Stephen Gilmore, Matthias Hölzl, Roberto Lucchi, Mirco Tribastone, and Gianluigi Zavattaro. *SENSORIA Process Calculi for Service-Oriented Computing*. In Ugo Montanari, Don Sannella, and Roberto Bruni, editors, *Second Symposium on Trustworthy Global Computing (TGC 2006), Lucca, Italy*, volume 4661 of *LNCS*, Italy, 2007. Springer.
 16. Martin Wirsing, Grit Denker, Carolyn Talcott, Andy Poggio, and Linda Briesemeister. A rewriting logic framework for soft constraints. *Electron. Notes Theor. Comput. Sci.*, 176(4):181–197, 2007.