

Modelling, Validating, and Ranking of Secure Service Compositions

Achim D. Brucker¹, Bo Zhou², Francesco Malmignati³, Qi Shi², and Madjid Merabti⁴

¹ The University of Sheffield, Sheffield, United Kingdom. E-mail: a.brucker@sheffield.ac.uk

² Liverpool John Moores University, Liverpool, United Kingdom. E-mail: {b.zhou, q.shi}@ljmu.ac.uk

³ Selex ES S.p.A, A Finmeccanica Company, Italy. E-mail: francesco.malmignati@guests.selex-es.com

⁴ College of Sciences, University of Sharjah, Sharjah, UAE. E-mail: mmerabti@sharjah.ac.ae

SUMMARY

In the world of large-scale applications, software-as-a-service (SaaS) in general and use of micro-services, in particular, is bringing service-oriented architectures (SOA) to a new level: systems in general and systems that interact with human users (e.g., socio-technical systems) in particular are built by composing micro-services that are developed independently and operated by different parties. At the same time, SaaS applications are used more and more widely by enterprises as well as public services for providing critical services, including those processing security or privacy of relevant data. Therefore providing secure and reliable service compositions is increasingly needed to ensure the success of SaaS solutions. Building such service compositions securely, is still an unsolved problem.

In this paper, we present a framework for modelling, validating, and ranking secure service compositions that integrate both automated services as well as services that interact with humans. As a unique feature, our approach for ranking services integrates *validated properties* (e.g., based on the result of formally analysing the source code of a service implementation) as well as *contractual properties* that are part of the service-level-agreement and, thus, not necessarily ensured on a technical level. Copyright © 0000 John Wiley & Sons, Ltd.

Received ...

KEY WORDS: Service design, human-centred service compositions, service modelling, service deployment, service ranking, secure service composition, service availability, SecureBPMN

1. INTRODUCTION

Enterprises need to flexibly adapt to new processes and react on changes on the market quickly. In today's interconnected world, this also impacts enterprise IT systems: they also need to be flexible to support the business needs (see [30] for an overview of flexible enterprise IT approaches). At the same time more large-scale applications for enterprises, the public-sector, as well as for consumers are built using compositions of micro-services and are delivered as software-as-a-service (SaaS). The underlying paradigm (as well as technologies such as WSDL [23] or REST [29]) is not new: it was already promoted a decade ago in the form of the service-oriented architecture (SOA) [27, 39]. Still, the problem of how to provide service-oriented systems that are *secure by default* is unsolved: the lack of security is a main factor that hinders cloud adoption [8].

Security is often considered to be technical topic that is addressed by specialists. While this is certainly true for low-level technical security decisions (e.g., selecting a specific encryption scheme), there are many security aspects (such as access control or compliance) that need to be addressed right from the requirements elicitation phase. Moreover, these business-level security requirements are pre-requisite to many technical security decisions.



Even if security requirements are captured early on, there is still the problem of tracing their actual implementation—which might require their refinement or even the “translation” of business concepts into technical concepts. While this is true for all systems, it is a particular challenge for service-oriented systems as, usually, service developers that work with service compositions have only limited influence on the security of the composed services. Thus, the compositions need to select the most suitable ones on offer.

In our approach, a service developer constructs a service composition plan for a system that can contain both automated services as well as human-centred services. While often, this composition plan is driven by the functional system requirements, our approach supports the specification of security requirements as first class citizens. This allows to discuss these important non-functional requirements with customers early in the design-phase and, thus, helps to realise a development methodology that supports “security-by-design.” After searching for suitable services in a marketplace, the abstract composition plan will be associated with concrete services for each task in the plan. The selection of services guarantees the fulfilment of both the functional as well as the security requirements of the system.

An important part of building secure service compositions is the selection of the most appropriate—in terms of security as well as functionality—services for building the actual service composition. It inevitably involves the quantification and ranking of services, according to their security levels. The three most important pre-requisites for quantifying and ranking services are:

1. a model or specification of both the security properties that a service composition as a whole as well as each individual service needs to fulfil, and the security guarantees offered by the services.
2. an understanding of to what extent the security guarantees of a service can be trusted.
3. a ranking algorithm that can cope with uncertainties or weak security guarantees and still ensures that the service composition provides the required functionality and the needed level of security.

Our *contributions* in this paper address these requirements by developing an *integrated development process and framework* that supports security properties, as first class citizen; right from the beginning of the service composition process, in which the service composition is modelled for binding with the required security properties. Secondly, this model is formally analysed to ensure that the composition provides the actual security requirements based on the minimal guarantees provided by the services being composed. Thirdly, we employ a ranking approach that allows to select the most suitable services by considering both formally verified security properties as well as informally stated security properties that are guaranteed by contractual or legal frameworks.

Our implementation is integrated into the Aniketos framework [13]. The framework, including the implementation presented in this paper, is available as Free Software (<https://github.com/AniketosEU>).

This paper extends our previous works [12, 13, 16, 17, 24, 56] in several key aspects: first, the set of properties that can be analysed both on the implementation level as well as on the actual service compositions are extended, e. g., to support the analysis of cryptographic properties. Second, formal analyses that yield in a binary “secure” or “inconclusive” result are integrated with quantitative ranking approaches. Finally, the isolated modelling and analysis approaches are, for the first time, integrated into a uniform, tool-supported, process that supports the whole life-cycle of modelling and implementing secure systems based on a SOA.

The rest of the paper is organised as follows. The next section (Sect. 2) provides an overview of existing SOA frameworks and explains how to use the BPMN modelling tool to construct service composition. We follow up on this by explaining in Sect. 3 how to model secure services and secure service compositions. In Sect. 4, we present techniques for formally validating security properties of atomic services as well as service compositions. We introduce a ranking and quantification approach that takes this uncertainty into account in Sect. 5. We discuss the framework in which our solution is integrated in Sect. 6. In Sect. 7, we briefly present a case study and, finally, we discuss related work (Sect. 8) and draw conclusions (Sect. 9).

2. BACKGROUND: SOA AND BPMN

In this section, we introduce SOA, its security requirements as well as BPMN as a solution for describing systems that are built using service compositions *and* support both automated services (tasks) as well as human-centred services (tasks).

2.1. Service-Oriented Architecture and Its Security

A *service* is a unit that provides a certain functionality. The SOA allows users to reuse existing services depending on their requirements. Therefore services can be composed to form a larger application in an *ad hoc* manner. SOA platforms provide a foundation for modelling, planning, searching for and composing services. They specify the architectures required, as well as providing tools and support for service composition standards.

To facilitate service composition across different platforms, service modelling languages are used to describe a) the business requirements of a system and b) system resources. By expressing behaviour processes and system organisation in agreed formats, not only the services can be easily understood and composed, but also the compositions can be validated against desired criteria and modified to suit required changes in operation.

Security in SOA becomes a big challenge due to the lack of common ground. One service developed with good faith in its security may not be necessarily good enough for another to use. For instance data access is a security issue that concerns most information systems. It is one of the main objectives while deploying secure services. Weak access control can cause severe consequences such as information leakage or data integrity issues. The situation gets more complicated in SOA as individual services from different domains may apply data access control in different—and often incompatible—ways.

Enterprise servers such as Glassfish do offer security parameterisations, but these are typically domain or platform-specific [21]. Subsequent standards have been proposed to augment the basic description of WSDL, to add semantic, behavioural, and to a limited extent, authentication and security data [3]. Other such property-based extensions, including Unified Services Description Language (USDL) [41], consist of standards that target trust and security, to bridge the previously-identified vendor divide.

2.2. Using BPMN to Construct Service Compositions

In process-oriented approaches, a service composition is often described using BPMN [46]. The modelling in BPMN is done by expressing business processes through business models. A BPMN model is a flowchart based diagram (see Fig. 1) that displays the basic structure and flow of activities and data within a business process.

2.2.1. Why BPMN. In our approach, we are using BPMN for modelling service compositions in the context of, e.g., business process-driven systems and socio-technical systems, i. e., systems that comprise machine-to-machine as well as human-to-machine interactions. The ability of BPMN to model both human as well as service tasks was one of the main reason for choosing BPMN over one of the many alternatives, including BPEL [47]. Moreover, BPMN is equipped with a standardised graphical notation that allows the visual modelling notation that is easy to understand and already known by many business experts. Finally, BPMN is executable and widely supported by multiple modelling and execution environments including Free Software implementations such as Activiti BPMN or SAP Netweaver BPMN. Thus, there is no need to translate BPMN to BPEL.

The approach presented in this paper is fully supported by a prototype developed on top of Activiti BPMN (based on BPMN 1.0) and, moreover, selected parts of the approach such as the secure modelling of BPMN as well as the validation of selected security properties are also available as part of an implementation based on SAP Netweaver BPMN (supporting a subset of BPMN 2.0).

2.2.2. Developing Service-Based Systems Using BPMN. From a high-level perspective, the development of a service-based system is divided into two phases:

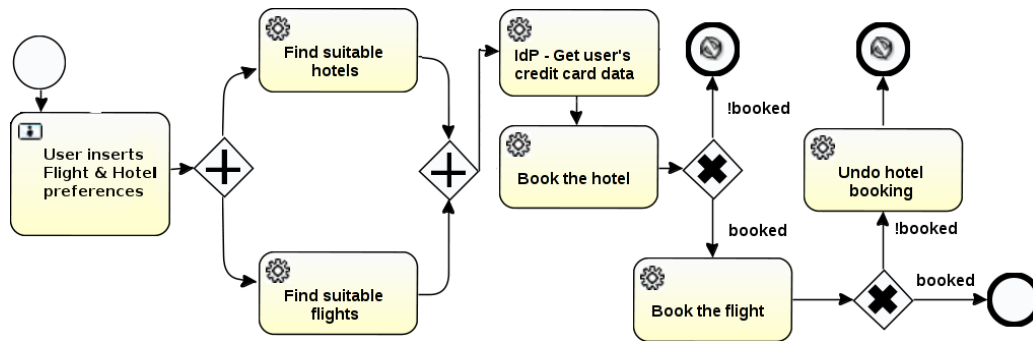


Figure 1. A service composition of a service for booking travels.

1. In the *design phase*, a service developer (together with domain or business experts) designs the process model, representing the composition of automated and human-centred service.
2. In the *deployment phase*, the process model is deployed in a business process execution engine, which can act as a service orchestrator as well as manage and configure other parts of the runtime infrastructure, e. g., enforcing access control.

This high-level view does not include several other tasks involved in system development, e. g., the implementation of actual services and design of user interface.

Fig. 1 shows a very simple BPMN diagram modelling a service composition that provides a travel booking service to customers. First, the customer enters his/her flight and hotel preferences into the system (such user interactions are modelled by user tasks in BPMN). Next, two web services (modelled as service tasks) are executed and connected via parallel gateways. These web services can be operated by different service providers and, in our example, provide functionalities for finding suitable hotel and flight information respectively. Here the parallel gateways ensure that the service which queries the customer's credit card data will only be executed if both the Find suitable hotels and Find suitable flights tasks terminated successfully. By using exclusive gateways the service developer is able to indicate that the Book the hotel task might fail. In case the booking fails (!booked), an error boundary event will be reached. Finally, the regular starting and ending points of the workflow are marked, respectively, by start and end events.

Modern systems need to fulfil a plethora of security requirements. In our simple example, to avoid fraud or price-fixing agreements, we could demand that the services for finding hotels and flights and the service doing the booking, are from different service providers. Moreover, only authenticated users will be allowed to authorise a booking. In addition, there might be also other requirements, e. g., only few service providers are "trustworthy" enough to handle the booking task.

3. MODELLING SECURE SERVICES

Modelling secure systems requires both the modelling of secure atomic services (i. e., automated services) and of secure service compositions that combine both services that interact with human users as well machine-to-machine communication.

3.1. Modelling Secure Atomic Services

We use ConSpec [4] for modelling atomic services as well as for specifying the (secure) input/output behaviour of composed services (i. e., composed services that are modelled as "black box"). Using ConSpec for our work is motivated by three reasons: 1. ConSpec was designed for specifying security properties, 2. it also supports the monitoring of security properties at runtime (e. g., see [7]), and 3. by using a language that is independent from the underlying service technologies, we can support different service technologies (e. g., RESTful services, WSDL-compliant services) at the

```

1  RULE ID ruleID
2  SCOPE <Session | Multisession>
3  SECURITY STATE
4    <bool | int | string> VarName1 = <Value1>
5      ⋮
6    <bool | int | string> VarNamen = <Valuen>
7
8  <BEFORE | AFTER> event1 PERFORM
9    Guard1,1 → Update1,1
10     ⋮
11   Guard1,m → Update1,m
12
13  ⋮
14
15  <BEFORE | AFTER> eventi PERFORM
16    Guardi,1 → Updatei,1
17     ⋮
18   Guardi,j → Updatei,j

```

Figure 2. The concrete syntax of ConSpec.

```

1  RULE ID Confidentiality_Booking
2  SCOPE Multisession
3  SECURITY STATE
4    string ServiceID = Hotel_Booking
5    string inputSuite = Basic256Sha256Rsa15
6    string inputSchema = symmetric
7    string inputAlgorithm = AES
8    int inputKeyLength = 256
9    string outputSuite = Basic256Sha256Rsa15
10   string outputSchema = symmetric
11   string outputAlgorithm = AES
12   int outputKeyLength = 256
13  BEFORE activity.start(string id, string type,
14                       int time, int date, string exec,
15                       string Output) PERFORM
16  ServiceID = id ∧ ...

```

Figure 3. A ConSpec specification requiring encryption.

same time. Our work can easily be adapted to other service specification languages that support security properties such as USDL [41], PROTUNE [11], or combinations of XACML [45] and WSDL [23].

ConSpec is a rule-based language (see Fig. 2 for its concrete syntax). The tag `RULE ID` simply defines the ID of the policy defined. The tag `SCOPE` specifies whether the rule is applied to one specific execution or to all executions of the service. The tag `SECURITY STATE` defines the global variables and their initial values. Then several events are checked `BEFORE` or `AFTER` the event occurrence. If an event occurred, we check guards one by one until we find the one which is satisfied. In this case certain security updates are performed. If no guards are fired for the event, then the further execution is not permitted (and some further security actions, like notifying the customer, are triggered). In case no security updates are needed but the further execution is allowed, there is a special action `SKIP`, which does not do anything but continues the execution. There is also a possibility of specifying an `ELSE` statement for the cases, when the further execution should be allowed even if no guards are fired (we omitted this option here for simplicity). A state can be seen simply as a specific assignment to the variables defined in the `SECURITY STATE` part. Naturally, the assignment set in the `SECURITY STATE` part defines the initial state. Actions are defined by the guarded events (specified between `<BEFORE | AFTER>` and `PERFORM`), i. e., by the name of the event (class and method), the set of its parameters and possible assignments for these parameters (in the case of `AFTER` the results of the event are also considered).

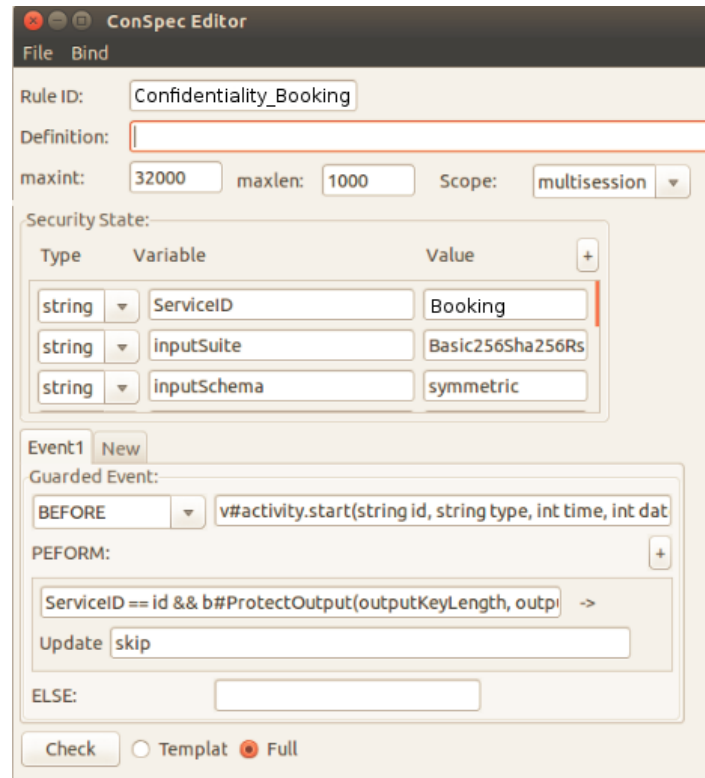


Figure 4. Confidentiality requirement in ConSpec editor.

Let us consider a candidate service for handling payments in our booking example (recall Fig. 1). Here, a natural requirement is the confidentiality of the credit card data, which can be achieved by using cryptography. Fig. 3 specifies this requirement using the concrete syntax of ConSpec.* This policy requires that both the input and the output of the service are encrypted with a specific cipher with a specific key length, that is part of the WS-Security cipher suite `Basic256Sha256Rs`. In more detail, we require the use of AES with a key length of (at least) 256 bits.

It is worth mentioning that the security specifications here are derived from user requirements. Therefore they can be fully customised by security experts. Instead of having to manually model everything, many of the requirements are directly derived from our Model Transformation Module, which is briefly mentioned in Sect. 6. ConSpec templates, which contain standardised security requirements such as encryption algorithms and authentication methods can also be created and enforced for each atomic services in the composition plan during the security validation phase as described in Sect. 4. It is of course less flexible and may not be very useful in certain scenarios.

In our framework, users (e.g., service developers) can use a user-friendly graphical editor for specifying ConSpec policies. For example, Fig. 4 shows how a confidentiality requirement (encryption) is configured and imposed on the Booking service.

3.2. Modelling Secure Service Compositions

Modelling security properties, as a first class citizen of a *service composition plan*, requires an integrated language for both security and functional requirements. We address this need with SecureBPMN, a meta-model-based [14] security language that integrates into BPMN. SecureBPMN extends BPMN 1.0 with means for specifying security properties. We based our work on BPMN 1.0 as at the point in time in which we started our work, BPMN 2.0 was not yet available. While already

*Our implementation is based on an XML ConSpec representation.

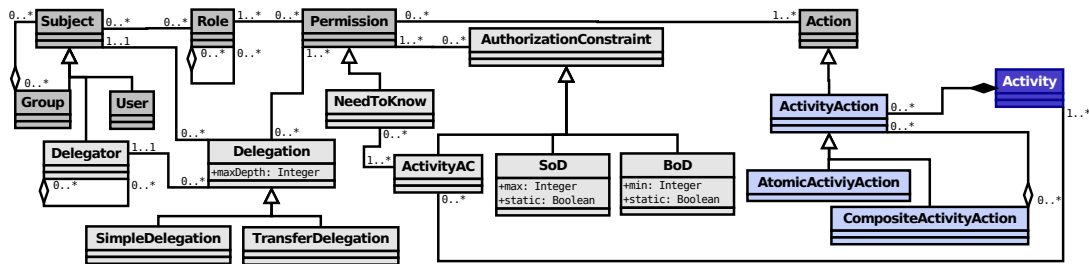


Figure 5. The SecureBPMN meta-model (simplified excerpt).

BPMN 1.0 allows for modelling simple security properties, neither BPMN 1.0 nor BPMN 2.0 are expressive enough to model the security needs of modern systems. For example, neither BPMN 1.0 nor BPMN 2.0 are able to express role-based access control constraints with dynamic and static separations of duty constrains. For a detailed discussion, we refer the reader to [12] (discussing SecureBPMN, extending BPMN 1.0) and [49] (discussing SecBPMN, extending BPMN 2.0).

Fig. 5 shows a simplified excerpt of the SecureBPMN meta-model that describes the domain-specific language for expressing the core security properties supported by SecureBPMN. The selection of security and compliance properties supported by SecureBPMN is based on discussions with various experts at SAP SE as well as the case studies conducted together with the industrial partners of the Aniketos project.

- *Role-based access control (RBAC)*: SecureBPMN provides a hierarchical role-based access control language supporting arbitrary constraints on the permissions. A **Subject** can be an individual **User** (i.e., a human user interacting with the system) or a **Group** of subjects. Subjects are mapped to a **Role** hierarchy. It is allowed in SecureBPMN to explicitly permit (**Permission**) the actions (**Action**) on the BPMN meta-classes **Activity**, **Process**, and **ItemAwareElement**. The latter is a class in the BPMN meta-model [46] from which, e.g., the BPMN **DataObject** is derived.
- *Permission-level separation and binding of duty*: SecureBPMN models separation of duty (**SoD**) and binding of duty (**BoD**) as sub-types of **AuthorizationConstraint**. SecureBPMN generalises the, usually binary, SoD and BoD constraints to n -ary constraints: an SoD constraint models that a **Subject** is not allowed to “use” more than \max permissions out of n ($\max < n$); BoD is generalised similarly. Finally, if a SoD (BoD) constraint already guaranteed by the RBAC configuration, it is called **static** SoD (BoD).
- *Delegation*: SecureBPMN supports delegation of tasks and, thus, the execution of services, with (**TransferDelegation**) and without (**SimpleDelegation**) transfers of the necessary access rights. The former only allows delegation of tasks to subjects that already possess the necessary rights. The latter allows delegation of tasks to arbitrary subjects that, then, can act on behalf of the original subject (**Delegator**). The number of delegations can be restricted by \maxDepth , e. g., a \maxDepth of zero forbids any delegation.
- *Need-to-know principle*: Confidentiality or a strict application of the need-to-know principle (**NeedToKnow**) is another important security property. In the context of service compositions this mainly refers to restrictions on the access to process variables or data objects (instances of the BPMN meta-class **ItemAwareElement**) and, thus, restricts the process of internal data-flow.

In our experience, these properties cover the most important needs for designing service compositions. For the more advanced features of SecureBPMN, such as the support for break-glass access control policies [18], history-resets for binding-of-duty with loops, or negotiable delegations, we refer the reader to [12]. To support service designers as well as security experts, we extended Activiti BPMN editor. This extended editor (see Fig. 7) supports the user-friendly modelling of security requirements.

```

1 <wsdl:definitions
2   xmlns:tns="http://booking.aniketos.eu/" ...>
3   <wsp:Policy ...>
4     <wsp:ExactlyOne>
5       <wsp>All>
6         <sp:SymmetricBinding>
7           <wsp:Policy>
8             ⋮
9             <sp:AlgorithmSuite>
10              <wsp:Policy>
11                <sp:Basic128Sha256Rsa15/>
12              </wsp:Policy>
13            </sp:AlgorithmSuite>
14          </wsp:Policy>
15        </sp:SymmetricBinding>
16      ⋮
17    </wsp:Policy>
18  </wsdl:definitions>

```

Listing 1: Excerpt of the WS-Policy for the service “Book the hotel.”

4. VALIDATING SECURITY PROPERTIES

After we have specified the security properties of atomic services as well as service compositions it would be nice, if we could verify that these properties hold for a specific instantiation (i. e., selection of services). While this is not possible for all properties, for many important properties it is achievable. In this section, we will present a verification approach to validate that services fulfil certain security properties. Together with contractual properties specified in the service level agreements, the result of the validation serves as input to the ranking and quantification process.

4.1. Validating Atomic Services

Atomic services are realised by implementing the business logic as computer program which is usually offered as a service (e. g., by deploying it as WSDL compliant web services). Recall our booking service (Fig. 1), assume that we want to implement this service as a WSDL compliant web service. The implementation of the **Book the hotel** service should fulfil, among others, the following two security properties:

1. as the information sent to the service is confidential (e. g., the travel destination or the credit card data), the service *shall only accept encrypted data as input* (as specified in Fig. 3) and
2. to minimise the attack surface as well as ensure compliance to the Payment Card Industry (PCI) Data Security Standard (<https://www.pcisecuritystandards.org/>), the Card Verification Number Scheme (“CVS Code”) *shall not be stored* (e. g., in a database).

To validate these properties, two artefacts need to be checked: for the first property, we need to analyse the WS-Policy configuration (see Sect. 4.1.1); for the second property, we need to analyse the actual source of the service implementation (see Sect. 4.1.2).

4.1.1. Validating Service Configurations. Listing 1 shows a simplified version of the WS-Policy [52] (respectively, WS-Security) specification for the **Book the hotel** service. This policy specifies that both input and output of the web service are encrypted using the algorithm suite `Basic128Sha256Rsa15` (line 11) during transmission.

Now, recall the requirements specified in ConSpec (Fig. 3). On the first glance, the WS-Policy specification seems to comply with the ConSpec specification. However this is not true: the ConSpec specification requires to use encryption keys with length of 256 bits while the WS-Policy only uses keys with length of 128 bits. Thus, such an implementation of the **Book the hotel** service *does not* fulfil our security requirements.

To detect this kind of configuration problems, we implemented a service that is able to check the followings:

- the https configuration of a server: we check the validity of the server's certificate chain with respect to a user configurable list of trusted root CAs and supported ciphers.
- the WS-Policy configuration for a WSDL-compliant web service. Here, we check in detail the SOAP messages that represent the parameters and return values of the service calls.
- the framework configuration for a RESTful service in a more complex situation. First and foremost, we check the https configuration of the framework. Moreover, if the framework used for implementing the RESTful service supports additional means for configuring security, we check this configuration as well.
- the frameworks, such as Apache CXF, allow to configure authentication and authorisation in a declarative way. Together with runtime information such as the user-role mapping, we check if the authentication and authorisation are configured according to the requirements expressed in ConSpec.

As the configuration is analysed and compared with the security requirements (i. e., the ConSpec specification), there are two checking modes:

1. *strict*: we require that the actual configuration matches exactly the requirements. For example, the key lengths must be exactly the same.
2. *relaxed*: we check that the actual configuration is at least as secure as specified in the requirements. For example, a service configuration using RSA with key length of 512 bits satisfies the requirement of using RSA with key length of 256 bits.

The *relaxed* checking is only used for properties, such as key length, that clearly provide a higher level of security. In particular, we do not allow *relaxed* checking on the actual cipher algorithms (e. g., RSA, AES), nor their mode (e. g., CBC, EBC), as the selection of such important properties should be a careful decision made by security experts.

As default we still recommend the *relaxed* checking mode as it results in a larger set of candidate services. Thus it allows for greater flexibility during service composition while still ensuring the security requirements. As the implementation of these checks is straight-forward, we omit them due to space limitation.

4.1.2. Validating Service Implementations. Only a small fraction of security properties can be implemented by providing an appropriate configuration. Most security measures are part of the service implementation, i. e., they are part of the computer program that implements the service. Listing 2 sketches a simplified implementation of the credit card validation code for the web service *Book the hotel* and lets have a closer look: this service handles credit card data and, thus, needs to comply to the PCI standard—even though the credit card is usually not charged when reserving a room, the card data is validated to ensure the possibility to charge the card in case of a late cancellation. The PCI standard states explicitly that the CVS of the card shall not be retained on the system. We can state this property in ConSpec as follows:

```

1  RULE ID CvsNotRetained
2  SECURITY STATE
3  BEFORE retain(...) PERFORM
4  checkForArgumentCvs -> {skip>

```

where `retain(...)` is a virtual function that captures all method calls that potentially will retain the data, e. g., write access to a database or the file system and `checkArgumentForCvs` is a predication that checks for any arguments containing the CVS. This data flow depended predication uses a combination of a naming heuristic (e. g., if a variable's name contains CVS, we assume it stores a CVS) and a data flow analysis (e. g., for known APIs that provide access to the CVS such as the result of the *Get user's credit card data* service).

```

1 //read input parameters
2 String arrivalDate = req.getParameter("arrival");
3 String guestName = req.getParameter("guestName");
4 :
5 // read credit card data from IdP
6 String ccHolder = service.getCcData("ccHolder");
7 String ccNumber = service.getCcData("ccNumber");
8 String ccCVS = service.getCcData("ccCVS");
9 String ccValidity = service.getCcData("ccVal");
10 :
11 // check credit card
12 log.write("Check_Credit_Card_for_guest_"+guestName);
13 if(validateCC(ccHolder,ccNumber,ccCVS)){
14     log.write("CC_validation_("+ccHolder+"/"
15             +ccNumber+")_successful.");
16 }else{
17     log.write("CC_validation_("+ccHolder+"/"
18             +ccNumber+"/"+ccCVS+")_failed.");
19 }

```

Listing 2: Simplified excerpt from the “Book the hotel” service.

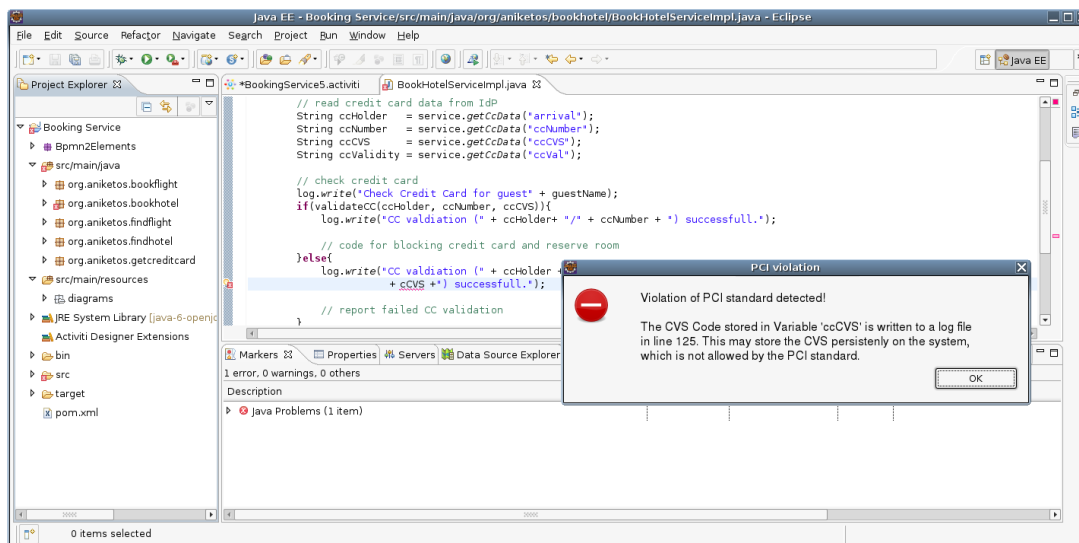


Figure 6. An example of an information disclosure violating the PCI standard.

We use static source code analysers to check the data flow within an implementation to ensure that the actual implementation does not violate the requirements. For this example we need to check that there is no execution path of the service implementation that accesses the CVS from the IdP and calls functions that can retain the data.

When we execute this check on our example service implementation (see Listing 2), we have been notified that the programmer implementing this service made a mistake: the CVS stored in the variable `ccCVS` (line 7) is written to a log file (line 18) and, thus, is retained in the file system or a database. Fig. 6 illustrates this notification. The Activiti Designer in Aniketos automatically switches to the Java perspective of Eclipse and highlights the source code that violates the security or compliance requirement specified in the context of the service composition.

We also check for common programming related security vulnerabilities such as SQL Injection and log file forging. The service designer does not need to specify these rather basic properties as we assume that they need to be fulfilled by all services. In our example, such a vulnerability will be reported as the content of the variable `guestName` (line 3), which can be influenced by an attacker, is written without any checks into the log file (line 12).

In the prototype we use our own static code analysis tool that is based on Wala (<http://wala.sf.net>). As an alternative, we can also generate configurations for a commercially available static code analysis tool.

4.2. Validating Security Compositions

Composing secure services does not, automatically, result in a secure service composition. On the one hand, there are properties, such as separation of duties, that inherently cannot be expressed on the level of an atomic service and, on the other hand, secure services can be used wrongly (e. g., using insecure configurations or sending confidential data to a public service).

To address these issues, we use an analysis method inspired by the work of [6]. We extended their work significantly to support n -ary SoD (BoD) constraints as well as constraints on the level of constrained permission (instead the task-level). As [6], we use the AVANTSSAR tool suite (www.avantssar.eu) as back-end for our formal analysis. Consequently, we translate the service composition plan and its security requirements to ASLan [6], i. e., the input language of the AVANTSSAR tool suite. The choice of ASLan is based on two reasons: 1) the experiments carried out by [6] show that ASLan is expressive enough to capture the requirements of security enriched service compositions and 2) the use of the same tools allows for developing a common verification back-end for our SecureBPMN-based approach as well as the approach developed by [6]. In fact, we could show that the analysis can be provided as a cloud-based service thus can be used by both modelling approaches [24].

Adding constraints such as SoD or BoD to a system that is already restricted by RBAC results in questions like the following: “Is the SoD constraint already guaranteed by the RBAC configuration?” Let us consider an RBAC configuration in which task t_1 can only be executed by members of the role r_1 and task t_2 can only be executed by members of the roles r_2 . Furthermore, let us assume that no user is assigned to both roles (i.e., no user is a member of r_1 and r_2). Thus, an SoD constraint between t_1 and t_2 is enforced already by the RBAC configuration and, hence, we only need to check this constraint after changes to the RBAC configuration are made. We call this a *static separation of duty* constraint. In contrast, let us consider an RBAC configuration in which tasks t_1 and t_2 can both be executed by members of the role r_1 . In this situation, an SoD needs to be checked, at runtime, for each and every access control request. Thus, we call this a *dynamic separation of duty*.

While static separation of duty constraints do not need to be enforced at runtime and, thus, reduce the runtime costs, it requires to re-check the SoD constraints after each and every modification of the RBAC configuration (e. g., adding new roles, changing the role assignment of subjects). In contrast, dynamic separation of duty constraints require a runtime check for each access to a resource that is constrained by separation of duty. While this is more flexible, it requires additional resources and, thus, costs more at runtime. Moreover, additional security checks might result in delays for users and, thus, might reduce the usability of the system.

Assume, in our example (recall Fig. 1), we want to counterfeit fraud or price-fixing agreements. Therefore, we require that the services Find suitable flights and Book the flight are operated by different providers (and similarly, for the hotel booking). The actual RBAC configuration is inferred automatically from the information available in the service marketplace (i. e., the SLA).

Our formal analysis translates the security configuration (here, RBAC and SoD/BoD) as well as the security properties that should be verified into the formal language ASLan [6]. In our example, the result of this translation (only an excerpt) for the security looks as follows:

```

1 hc rbac_ac(Subject, Role, Task)
2   := CanDoAction(Subject, Role, Task)
3   :- user_to_role(Subject, Role), poto(Role, Task)
4 hc poto_T6 := poto(TravelAgency1, Find suit. flights)
5 hc poto_T7 := poto(TravelAgency1, Book the flight)

```

where `poto` facts describe which users or roles can execute/access a task.

The security goal is, in this case a SoD constraint between the services Find suitable flights and Book the flight:

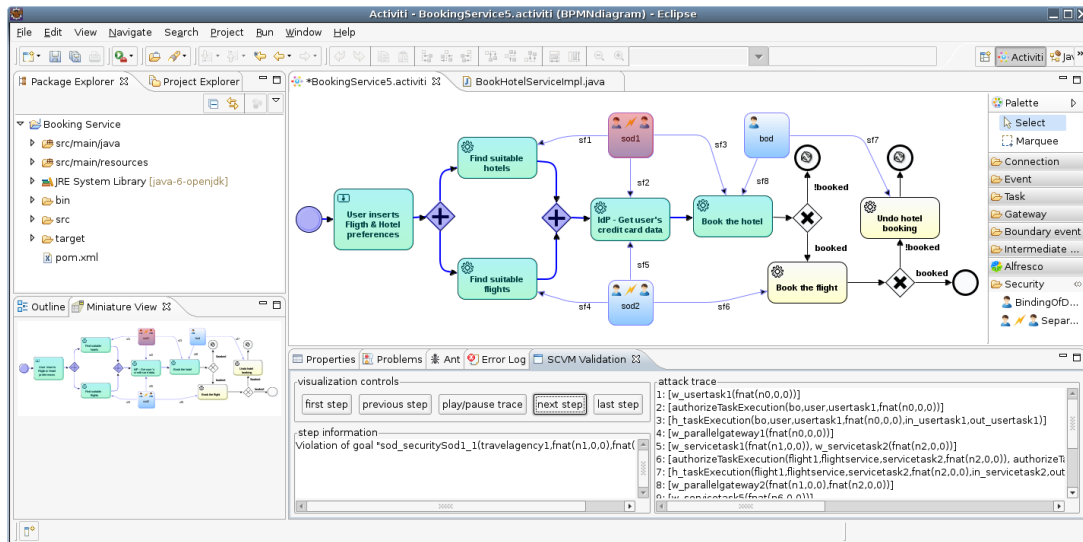


Figure 7. Security validation within the Activiti BPMN editor.

```

1  attack_state sod_securitySod1_1 (Subject0, Subject1,
2      Inst1, Inst2)
3  := executed(Subject0, task (Find suit. flights, Inst1)) .
4      executed(Subject1, task (Book the flight; Inst2))
5      &not (equal (Subject0, Subject1))

```

This configuration, obviously, violates the SoD constraint as the `TravelAgency1` can do both searching for flights and booking them. In this case, a dishonest travel agency could prefer flights with a higher bonus for the travel agency that are not necessarily the cheapest for the traveller. This is detected by our analysis, e. g., the verification modules returns the following “attack trace:”

```

1  1. [w_task1 (fnat (n0, 0, 0))]
2  2. [authorizeTaskExec (bo, user, task1, fnat (n0, 0, 0))]
3  3. [h_taskExec (bo, user, task1, fnat (n0, 0, 0),
4      in_task1, out_task1)]
5  4. [w_parallelgateway1 (fnat (n0, 0, 0))]
6  5. [w_servicetask1 (fnat (n1, 0, 0)),
7      w_servicetask2 (fnat (n2, 0, 0))]
8  6. [authorizeTaskExec (flight1, flightservice,
9      servicetask2, fnat (n2, 0, 0)),
10     authorizeTaskExec (travelagency1, travelagency,
11         servicetask1, fnat (n1, 0, 0))]
12     :
13 15. h_taskExec (travelagency1, travelagency,
14     servicetask9, fnat (n8, 0, 0),
15     in_servicetask9, out_servicetask9)

```

Of course, this textual representation is not well-suited to practitioners. Therefore, we developed a user-friendly visualisation of such an attack in terms of the high-level composition plan (i. e., on the level of the BPMN model). Fig. 7 shows how our prototype visualises such a violation to the service developer. The service developer is able to manually step through all necessary actions that a dishonest user would execute to violate the SoD constraint.

After such an analysis, the service developer needs to decide how to mitigate this risk. In general, there are several options, among them

- re-design the composition plan, to avoid the need for a particular separation of duty constraint,
- instruct the service composition framework to ensure the selection of different service providers, or

- enforce a dynamic separation of duty at runtime. For this, our prototype can generate configurations for XACML [45] based access control infrastructures.

The concrete mitigation plan depends on the actual use case.

5. SERVICE QUANTIFICATION AND RANKING

It is not uncommon, when composing service, that several service instances, offered from different service providers, fulfil the basic requirements. In this section, we will discuss this situation in more detail and present approach that supports service developers to select the “best” service according to their security needs.

5.1. The Role of SLA

The security property modelling and verification techniques allow the service consumer specify certain security properties that the service composition has to comply with. In practice, not all security properties are technically verifiable and some properties such as BoD and SoD are validated at design-time but not always enforced at runtime. Therefore we need to look at other sources that can provide security guarantees for web services.

Web services are normally made available together with a service-level agreement (SLA). A SLA is a guarantee that has to be accepted by service consumers before the service is used. A SLA can specify the properties of a service across different levels. For example, on business level it can describe what kind of functionality the service is offering and how the users will be charged (cost); on the technical level it may describe the number of shutdowns the service might experience each year (QoS).

Security can also be promised as part of the SLA. However its coverage is rather poor to date due to the lack of well defined semantics. The SLAs traditionally focus on the QoS metrics such as a bandwidth guarantee and backup strategy. Even when the security is mentioned, in practice it tends to be written in a natural language with fuzzy terms such as “*High*” or “*Good*.” Therefore it is very difficult for the service consumer to really understand the situation and compare the web services from the security perspective. Nevertheless as a legally bound document, SLAs are useful as a complement to technical verifications.

It is an interesting question to ask which security properties should be specified in the SLAs. As SLAs can be written in natural language, thus in theory it is possible to specify any security properties the service would like to offer. However, to make it meaningful and comparable, a proper schema must be defined first. Henning [33] was among the first trying to address the quantifiable security issue in SLAs by expressing and measuring the security of a service by associating it with performance related metrics. For example, a security requirement to “Restore backed up data” is measured by a quantifiable metric such as “Data restored 95% of times within a given response time.” The way the security has been expressed is rather subjective though, depending on the context of each enterprise, where the research was targeting. Therefore the process cannot be implemented automatically. Instead, it requires a close study of the enterprise’s configurations by security specialists. SecAg [31] [32] is another framework proposed to express security metrics in SLAs. SecAg extends the standard WS-Agreement [5] to provide necessary semantics for specifying security properties. For example, with the extensions it can specify which service level objective (SLO) is auditable and assign an access control list to the SLO. Based on the extensions, the author also proposed a risk-based approach for service matchmaking. Each SLO is assigned a weight w representing the risk that the SLO is not fulfilled. By calculating the weighted *Euclidean* distance of each SLA to the security requirements, using techniques such as a text similarity analyser, the SLA that is closest to the security requirements will be selected as the risk is at the minimum. With this solution though, there is a possibility that a SLA offering far better security may not be considered as the closest to the original requirements.

Despite of these efforts being made, the issue of measuring security of a service composition remains unsolved. In this section, we introduce the mechanism for quantifying and ranking

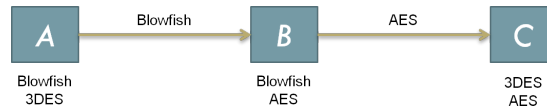


Figure 8. Example composition of three services.

service compositions, i. e., we support the service consumer in choosing, based on an automated recommendation, the most suitable service composition. This recommendation should be made based on the properties of the service composition as a whole, rather than just based on individual sub-services in the composition. The quantifying and ranking is used when the service consumers have to choose one from a number of available service compositions. It is particularly useful when the validation is not fine-grained, i. e., a large number of service compositions either pass or fail the validation altogether, which does not help to select the most suitable service composition.

As a starting point, we quantify and rank service compositions from three aspects, which are the three factors that are mostly considered by service consumers: encryption (security), availability (QoS), and cost (business). We focus on these three properties in this paper because not only they are normally mentioned in the SLAs, but also they are the properties that can be validated through different means. For example, the encryption algorithm is specified in SLA and can be validated by techniques explained in Sect. 4. Availability of a service can be easily recorded and calculated by examining the logs stored in the system. This work is implemented as a key part for the security composition planner module (see Sect. 6) in Anketos framework.

5.2. Encryption – The Weakest Link

There are some cases when the weakest link principle is particularly applicable to service composition. It states that when services are composed together, the security capability of the composite service is equal to what the weakest service or link offers. This security principle is applicable to many security properties and *encryption* is one of them. When encryption is applied to communications between services, the services may adopt different encryption algorithms or key lengths which give them different encryption strengths. To communicate with each other, the encryption strength of a service with an advanced encryption algorithm may be degraded by that of a service with a weak encryption scheme during the composition. Thus the composite service literally uses the weakest encryption strategy in part of its communications. For example, consider the case in Fig. 8 where service *A* supports encryption algorithms of Blowfish and 3DES, service *B* supports Blowfish and AES, and service *C* supports 3DES and AES. To communicate with each other, the link between service *A* and *B* is encrypted with Blowfish and the link between *B* and *C* is encrypted with AES. Therefore the overall strength of the composition, in terms of keeping communications confidential, is the weaker one between Blowfish and AES.

The weakest link principle is used to determine the security capacity of the service compositions. It should be noted however that the weakest link principle is not universally applicable. There are security cases where alterations to a service composition can be utilised to improve the security of a composite service to be greater than that of the weakest component. An example might be where a firewall service is used to shield an otherwise vulnerable service from outside attack. The use of the firewall mitigates the vulnerability exposed by the weaker service. And vice versa it may also apply in reverse: the introduction of a component may serve as an exacerbating factor that reduces the security of the overall composition to a degree beyond that posed by the service were it to act in isolation. This often results from interactions between incompatible security properties.

To simplify the issue, in this study we focus on the encryption. Therefore each link between services is checked, and the encryption strength of the composition is determined by the weakest link, i. e.,

$$E = \min_{i=1}^n E_i$$

Table I. Quantitative value of encryption algorithms.

Algorithm Name	Quantitative Value
Serpent	0.9
AES (Rijndael)	0.8
3DES	0.7
CAST128/256	0.6
Twofish	0.5
Blowfish	0.4
MARSH	0.3
Other encryptions	0.2
Codings	0.1
Plain text	0.0

where E is the encryption strength of the composition and E_i is the encryption strength for each link i in the composition. E_i is determined by the strongest algorithm supported by both services at each end of the link i .

The quantitative value (from 0.9 to 0 in our case), however is predetermined by expertise in advance based on Tab. I. As claimed in [37], the quantitatively ranking of encryption algorithms is possible but heavily depends on the metrics and target scenario. Tab. I is a guideline and rather used to demonstrate our ideas.

5.3. Availability

Availability is another aspect being used to compare services and it relates to QoS. Availability in this scenario means the available time ratio of a service. An unexpected service shutdown could cause severe damage to a service consumer's business and a service developer's reputation. Therefore seeking guarantee from the service provider about the service availability is one of the top priorities for service consumers, before they commit to use the service. The situation gets complicated in service composition because a composition's availability is decided by not only the technical specifications of the sub-services, but also by the structure of the composition.

Take the example of the travel booking service in Fig. 1 on page 4, where most of the services are placed in sequential order. That means if one of the sub-services is not available, the entire composition will stop. Therefore the availability of sequential tasks is the *product* of all the sub-services' availability values in percentage. However, the services Find suitable hotels and Find suitable flights are executed in parallel. It means these two services can be carried out separately. Nonetheless they still have to be both finished before the next task Get user's credit card data can be executed. Therefore for parallel tasks the availability value is the *minimum* among them. For services that are exclusive to each other, the availability of the composition depends on which service has been eventually used.

Tab. II shows the rules that we used for calculating the availability of composite services. Assume in Fig. 1 each service has the following availability value: Find suitable hotels: 0.99, Find suitable flights: 0.96, Get user's credit card data: 0.97, Book the hotel: 0.99, Book the flight: 0.98, and Undo hotel booking: 0.94. The *availability* value for a successful transaction will be calculated as:

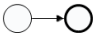


$$A = \min(0.99, 0.96) \times 0.97 \times 0.99 \times 0.98 = 0.90$$

where A represents availability of the composition.

5.4. Cost

Finally the last factor that also plays an important role in consumer's decision making is the *cost*. Higher security and QoS normally indicate a higher price, which must be within a consumer's budget. Comparing to *encryption* and *availability*, calculating the *cost* of a service composition is

Table II. Rules to calculate availability.

	Description	Calculation
	Sequence	$\prod_{i=1}^n A_i$
	Parallel	$\min(A_1, \dots, A_n)$
	Exclusive	A_i

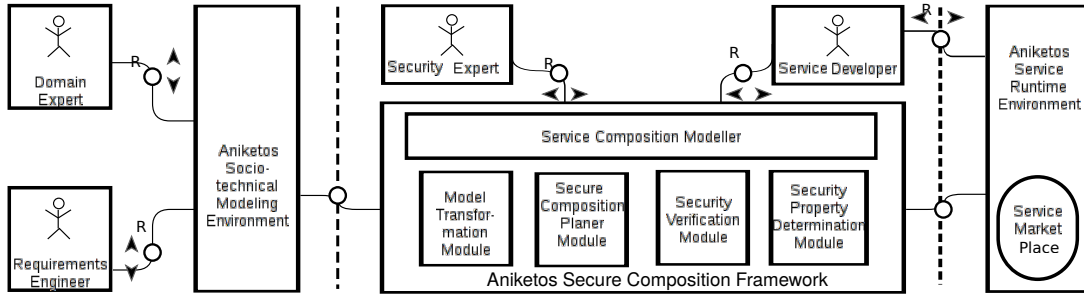


Figure 9. The Aniketos Secure Composition Framework.

more straightforward. It is the sum of all the employed atomic services' costs, i. e.:

$$C = \sum_{i=1}^n C_i$$

where C is the cost for the composition and C_i is the cost for atomic service i .

6. A SECURE COMPOSITION FRAMEWORK

Building secure composite services on top of a SOA is a challenging task. At *design-time* the service developer needs to select the optimal set of services that satisfies both the functional and security requirements put by the end user. At *runtime*, a service may become unavailable due to various reasons and has to be replaced automatically with an alternative service that, at least, offers the same security guarantees. In addition, the service developer also needs to decide if a given security property should be enforced statically or dynamically. While a static enforcement creates less overhead at runtime, it reduces the flexibility of service substitution or re-composition. In contrast, dynamic enforcement is usually more flexible but requires more system resources at runtime. Thus, a service designer needs also to consider economical aspects of realising security and compliance requirements.

To support the service developer in building flexible and secure services through compositions, we propose a *secure service composition framework* that addresses both the design-time and runtime secure service composition. We focus only on the technical parts of the design-time process, i. e., we exclude the requirements elicitation, as well as the service deployment and runtime adaptation parts.

Fig. 9 gives a high-level overview of the *Aniketos Secure Composition Framework* which is the design-time modelling and analysis part of the Aniketos platform. At the beginning, domain experts together with requirement engineers specify the high-level business process as well as the security requirements by using the *Aniketos Socio-technical Modelling Tool* [48]. It provides the opportunity to express security needs not just from technical, but also from social aspects (not discussed in this paper). From these semi-formal descriptions, the *model transformation module* automatically infers

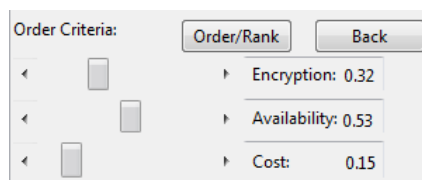


Figure 10. Set ranking criteria.

composition plans, which are presented in the BPMN format. These composition plans are coarse-grained. Thus, before these composition plans can be deployed in the *Aniketos Service Runtime Environment*, they will be refined by a service developer using the *Aniketos Secure Composition Framework*.

The *Aniketos Secure Composition Framework* provides an Eclipse-based environment (the *Service Composition Modeller*) to the service developer for refining the composition plans as well as checking their security properties. Specifically, the service developer can, among others, use the following component modules:

- *Model Transformation Module*: infers the draft composition plan from the requirement document expressed in the Aniketos Socio-technical Modelling Language [48].
- *Secure Composition Planner Module*: allows the service developer to semi-automatically select secure services for a given composition plan (see Sect. 5). To check that the compositions comply with the security requirements, this module uses the Security Verification Module and the Security Property Determination Module.
- *Security Verification Module*: provides formal validation and verification solutions for composed services and atomic services, as discussed in Sect. 4.
- *Security Property Determination Module*: provides a uniform interface for accessing security properties of services. Moreover, this module stores the verification status of security properties to avoid an unnecessary (expensive) re-verification.
- *Service Marketplace*: registers and stores the services for open access. The Secure Composition Planner Module selects services from the Service Marketplace.

The framework also includes a simple user interface providing prioritising options so that the service developer can specify the criteria used to rank the service compositions. As shown in Fig. 10, the service developer is able to choose how much weights he/she wants to put on each criterion of encryption, availability, and cost. Assume the developer sets the weights to 0.32, 0.53 and 0.15 respectively, the overall value V for each service composition will be:

$$V = 0.32 \times E + 0.53 \times A + 0.15 \times \frac{B - C}{B}$$

where E represents the value of encryption strength, A represents the value of availability, C represents cost, and B represents the consumer's budget. These values are calculated using the methods discussed in Sect. 5. Apparently higher values of E and A as well as a lower value of C will result in greater value of V . In this way the generated service compositions cannot only be security-wise verified by our SecureBPMN extensions, and also ranked easily based on the developer's other priorities.

7. A CASE STUDY

We implemented a prototype of our framework based on the Activiti BPMN tool suite (<http://www.activiti.org>). This prototype was applied to several industrial case studies within

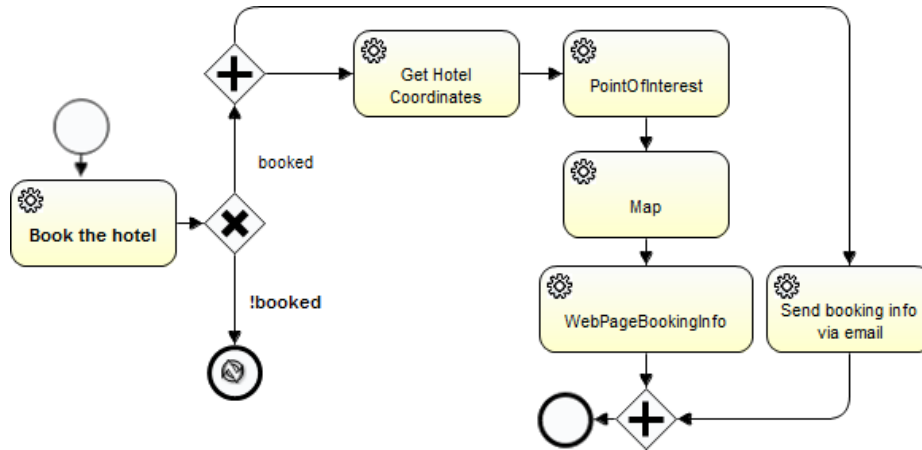


Figure 11. BPMN diagram of the booking hotel case study.

the Aniketos project. Additionally, we discussed our with domain experts from SAP and SAP's customers. In this section we illustrate a simple process of using the tool suite to book a hotel. Here, we focus on service tasks that can be executed automatically to demonstrate the idea of how services are composed to create new applications.

The case study presented in this section illustrated our overall approach. For the evaluation of our approach (see also the discussion in Sect. 9), we used three larger case studies from three different domains: air traffic management, public sector, and telecommunication. For details, we refer the reader elsewhere [1].

Our illustrative case study is as follows: to offer a user with helpful and smooth booking experience, a new hotel booking application needs to be composed by multiple services. Basically once a booking is made, we want to provide the user with some local information about the hotel, such as the point of interest, as well as an email confirmation. Specifically, providing the local information involves four services: 1) Get the hotel coordinates, 2) Retrieve point of interest around the hotel's coordinates, 3) Load the map around the hotel, and 4) Create a new web page to display these information. Together with the actual booking and email confirmation services, in total six services will work together to provide the new application. Each of the service in the application can be provided by more than one service providers, offering different security properties.

We model the system using BPMN. Actors are represented as roles that are assigned to tasks in the BPMN model. As shown in Fig. 11, six service tasks (*Book the hotel*, *Get hotel coordinates*, *Point of interest*, *Map*, *Web page booking info*, and *Send booking info via email*) are created in the BPMN diagram to represent the entire process from book a hotel, to display the booking information, and to send email notification to the user.

In the next step, security expert will specify security requirements following the steps explained in Sect. 3. One assumption here is that the atomic services will be registered first in our *Service Marketplace*, together with their SLAs. In this case study, two map services are registered for the *Map* task and both are discovered by the composition framework as shown in Fig. 12.

We registered two services for each of the six service tasks in the case study. Therefore in total it created 64 (2^6) possible service compositions (also called composition plan in our framework), which will not all pass the validation process described in Sect. 4. The remaining composition plans are ranked by the service developer based on user's preferences, as described in Sect. 5. Fig. 13 shows the results and this completes the design phase of the development of secure service composition.

Finally the chosen composition plan (normally the one ranked first) will be deployed to the *Aniketos Service Runtime Environment* (Activiti Engine-based) as shown in Fig. 14. Starting up the composite service will invoke the atomic services in turn and display the booking result as a

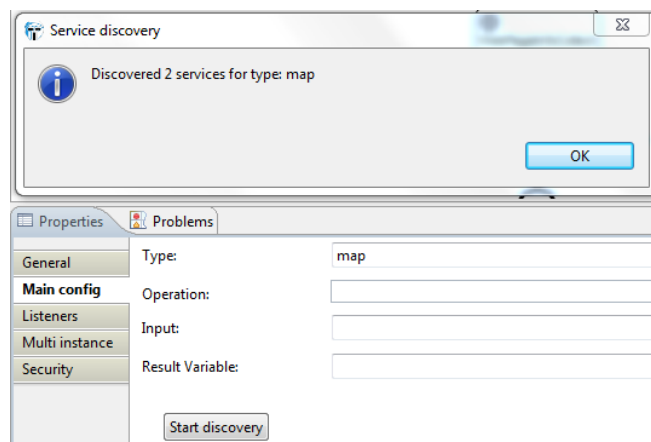


Figure 12. Two map services are discovered.

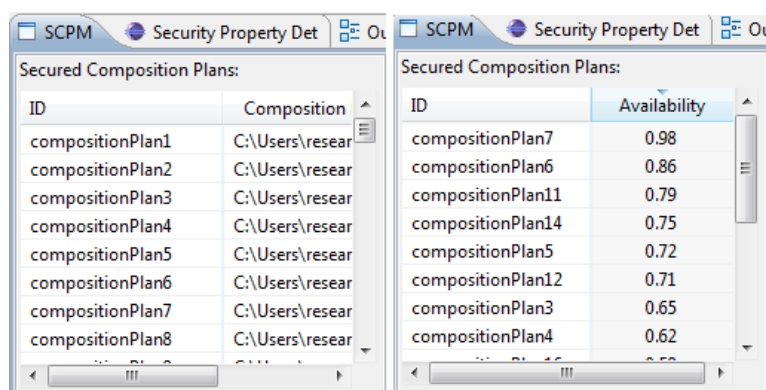


Figure 13. Creation and ranking of composition plans.

web page, as illustrated in Fig. 15. During runtime, the access control policies are enforced by an XACML-based infrastructure [17].

8. RELATED WORK

We see three areas of related work: 1. modelling of security requirements for process models, 2. analysing security properties of process models, and 3. determining security of composite services.

There is a large body of literature extending graphical modelling languages with means for specifying security or privacy requirements. One of the first approaches is SecureUML [40], which is conceptually very close to our BPMN extension. SecureUML is a meta-model-based extension of UML that allows for specifying RBAC-requirements for UML class models and state charts. There are also various techniques for analysing SecureUML models, e. g., [10] or [15]. While based on the same motivation, UMLsec [38] is not defined using a meta-model. Instead, the security specifications are written, in an ad-hoc manner, in UML profiles. Similar to UMLsec, [44] presents an attribute-based approach (i. e., the conceptual equivalent of UML profiles) of specifying security constraints in BPMN 2.0 models. Inspired by these works, there are several approaches extending BPMN 2.0 with security specifications, e. g., [22, 49]. While they provide, on the one hand, further security properties that are not supported by SecureBPMN, they all have in common that the access control specifications are very coarse-grained (only supporting simple RBAC models). In contrast,

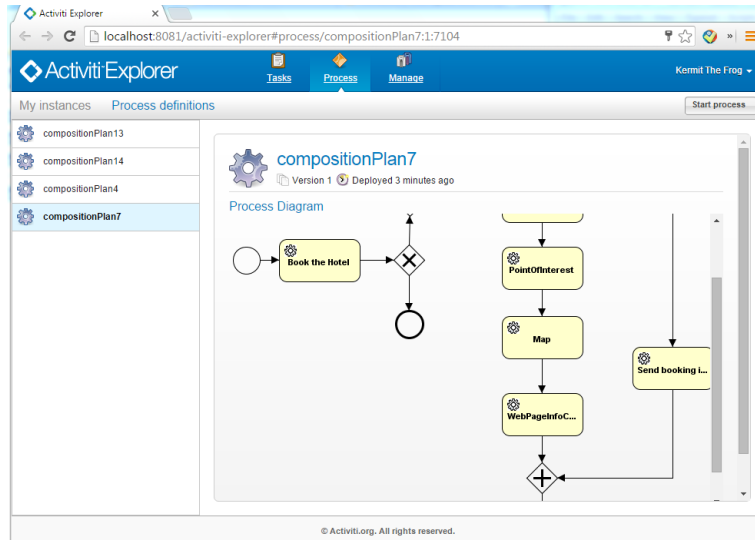


Figure 14. Composition deployed to the runtime environment.

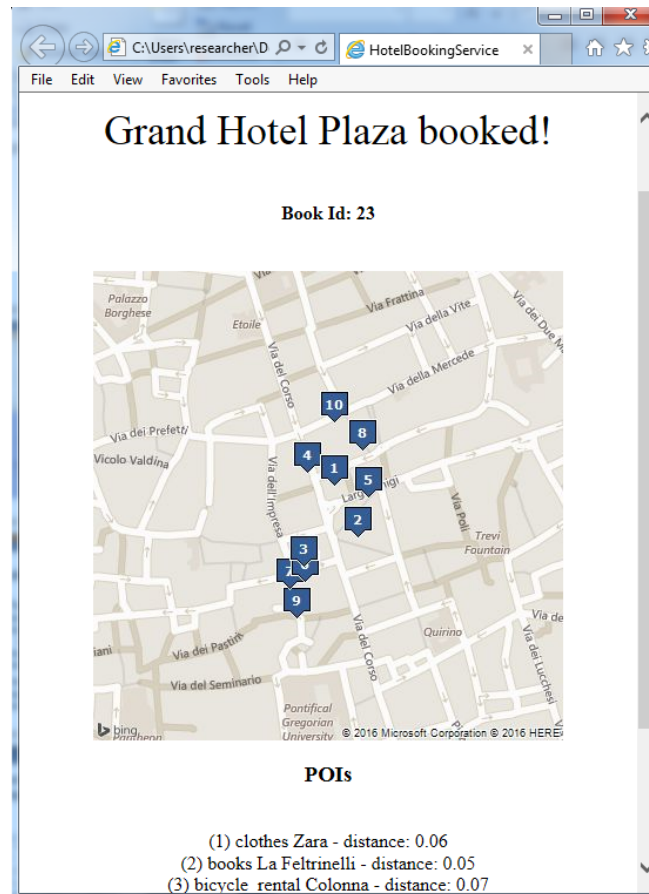


Figure 15. Running service-based application of the case study.

our approach allows the fine-grained specification of security requirements for single tasks or data objects.

With respect to the validation of security requirements on the business process level, the closed related work is the work of [53] and [6] that both support the checking if an access control specification enforcing binary separation of duty and binding of duty constraints. Apart from security properties, there is also a strong need for checking the consistency of the business process itself, e. g., the absence of deadlocks. There are several works that concentrate on this kind of process with internal consistency validation, e. g., [25] and [2]. Moreover, there are several approaches for analysing access control constraints over UML models, e. g., [51], [15], and [38]. These approaches are limited to simple access control models, as the UML models are usually quite distant from business process descriptions comprising high level security and compliance goals.

Last but not least, determining the properties of a composite service based on its atomic services is another area that attracts attentions from the research community. To achieve this, the first step is to quantify web services. In the past the focus was on raking web services based on just their QoS metrics and trying to find the best match. Paper [50] ranks web services under multi-criteria matching. It targets at accurate web service selection and assigns a dominance score to each advertised web service. [43] defines a business-focused ontology to enable semantic matchmaking in open cloud markets. Paper [35] proposed the concept of Quality of Security Service. It treats the security as part of QoS requirements. The author argues that security requirements such as the strength of a cryptographic algorithm, the length of a cryptographic key, security functions, confidence of policy-enforcement and the robustness of an authentication mechanism would all be specified and measured as the quality of security services. Paper [19] proposed an Analytic Hierarchy Process (AHP) based framework for web service quality evaluation. It uses a quality meta-model to format SLAs and assigns weights to different quality characteristics based on their importance. Similarly [20] uses a Singular Value Decomposition (SVD) based technique, and a user assisted weighting system to find higher order correlations among web services. With respect to the determination of properties of composite services, paper [36] focuses on the QoS values of service composition. It takes the structure of the services into account, in a similar way as we determine the availability of the service compositions. Based on a process sequence such as *loop*, *and*, *or*, the QoS values are calculated according to predefined rules. Elshaafi et al. [26] use a similar method but the focus was on the trustworthiness of service composition. The authors argue that trustworthiness of a service composition is a combination of properties such as reputation, reliability and availability etc. These properties are one step closer towards general security issues and the authors are also from the Aniketos framework development team. Zhou et al. [56] propose a classification method that abstracts and quantifies service compositions based on five key security aspects: confidentiality, integrity, availability, accountability and non-repudiation. There are also other works that focus on security properties of system-of-systems such as [55] and [54]. Comparing to these works, our approach concentrates on the most objective and justifiable properties in encryption, availability and cost, which represents security, QoS and business respectively. Our solution also gives flexibility to the end users so that they can decide how to prioritise these properties in service compositions.

9. CONCLUSIONS AND LESSONS LEARNED

We presented a practical approach for developing service-oriented systems. Our approach supports certain security properties following the “secure-by-design” paradigm. Our approach focuses on the most important “high-level” security properties to allow non-security experts (e.g., business analysts) to consider security right from the beginning. As such, it does not replace traditional secure development processes and techniques [9, 28, 34, 42] that address architectural and implementation security aspects.

Besides the presented illustrative case study and three large case studies in the context of the Aniketos project[1], we discussed our approach with various experts at SAP. Overall, these experiences show that our approach is applicable to a wide range of application domains.

Our approach provides a seamless integration of technical security properties that can be formally verified with (informal) security and other requirements that are specified in SLAs. This was rated as a unique and very powerful feature of our framework. The experts at SAP suggested to extend this approach even one step further to include post-hoc checks: as in many systems properties such as separation of duty are not enforced at runtime, they need to be checked—during audits—by analysing the log files.

Moreover, our interview partners for the Aniketos project liked that security properties can be modelled by non-security experts together with the service composition. While it is understood that all security requirements need to be reviewed and extended by security experts, offering non-security experts the possibility to initially model their security requirements was seen as a competitive advantage.

Our evaluation showed that the supported security properties are sufficient for most modelling needs. Still some case studies raised the need for various notions of confidentiality. Confidentiality, in terms of requiring encrypted communications between the different services (tasks) is an important requirement. Choosing the correct encryption techniques (in fact, on a technical level, we need to ensure that data is only communicated over authenticated and secure channels) requires a multitude of technical decisions (e. g., encryption algorithms, and length of cryptographic keys).

Finally, our evaluation showed our formal analysis is usually able to validate security or compliance properties within less than 20 seconds. While this is fast enough for the (interactive) design of service compositions, it is too slow for automatic service re-composition at runtime. Therefore, the efficient caching, which needs to ensure the authenticity and validity of validation results is of outermost importance.

REFERENCES

- [1] Deliverable 6.4: Final report on aniketos applied to industrial case studies. Tech. rep., Aniketos (2014). URL <http://www.aniketos.eu/sites/default/files/downloads/Aniketos%20D6.4%20-%20Final%20report%20on%20Aniketos%20%20applied%20to%20industrial%20case%20studies.pdf>
- [2] van der Aalst, W.M.P., Dumas, M., Gottschalk, F., ter Hofstede, A.H.M., Rosa, M.L., Mendling, J.: Correctness-preserving configuration of business process models. In: Fiadeiro, J.L., Inverardi, P. (eds.) *FASE, LNCS*, vol. 4961, pp. 46–61. Springer (2008)
- [3] Akkiraju, I.R., et al.: Web service semantics – WSDL-S (2005)
- [4] Aktug, I., Naliuka, K.: Conspec - A formal language for policy specification. *ENTCS* **197**(1), 45–58 (2008)
- [5] Andrieux, A., Czajkowski, K., Dan, A., Keahey, K., Ludwig, H., Pruyne, J., Rofrano, J., Tuecke, S., Xu, M.: Web services agreement specification (WS-Agreement). Tech. rep., Open Grid Forum (2007)
- [6] Arsas, W., Compagna, L., Pellegrino, G., Ponta, S.E.: Security validation of business processes via model-checking. In: Erlingsson, Ú., Wieringa, R., Zannone, N. (eds.) *ESSoS, LNCS*, vol. 6542, pp. 29–42. Springer (2011)
- [7] Asim, M., Yautsiukhin, A., Brucker, A.D., Lempereur, B., Shi, Q.: Security policy monitoring of composite services. In: Brucker, A.D., Dalpiaz, F., Giorgini, P., Meland, P.H., Rios, E. (eds.) *Secure and Trustworthy Service Composition: The Aniketos Approach*, no. 8900 in *LNCS: State of the Art Surveys*, pp. 192–202. Springer (2014)
- [8] Autotask Corporation: Metrics that matter (2014). <http://www.autotask.com/lp/metrics-that-matter-2014/> (2014)
- [9] Bachmann, R., Brucker, A.D.: Developing secure software: A holistic approach to security testing. *Datenschutz und Datensicherheit (DuD)* **38**(4), 257–261 (2014). doi: 10.1007/s11623-014-0102-0
- [10] Basin, D., Clavel, M., Doser, J., Egea, M.: Automated analysis of security-design models. *Information and Software Technology* **51**(5), 815–831 (2009)

- [11] Bonatti, P.A., Coi, J.L.D., Olmedilla, D., Sauro, L.: A rule-based trust negotiation system. *IEEE Trans. Knowl. Data Eng.* **22**(11), 1507–1520 (2010)
- [12] Brucker, A.D.: Integrating security aspects into business process models. *it* **55**(6), 239–246 (2013)
- [13] Brucker, A.D., Dalpiaz, F., Giorgini, P., Meland, P.H., Rios, E. (eds.): *Secure and Trustworthy Service Composition: The Aniketos Approach*. No. 8900 in LNCS. Springer (2014)
- [14] Brucker, A.D., Doser, J.: Metamodel-based uml notations for domain-specific languages. In: Favre, J.M., Gasevic, D., Lämmel, R., Winter, A. (eds.) *ATEM* (2007)
- [15] Brucker, A.D., Doser, J., Wolff, B.: A model transformation semantics and analysis methodology for SecureUML. In: Nierstrasz, O., Whittle, J., Harel, D., Reggio, G. (eds.) *MoDELS*, no. 4199 in LNCS, pp. 306–320. Springer (2006)
- [16] Brucker, A.D., Hang, I.: Secure and compliant implementation of business process-driven systems. In: Rosa, M.L., Soffer, P. (eds.) *SBP, LNBIP*, vol. 132, pp. 662–674. Springer (2012)
- [17] Brucker, A.D., Hang, I., Lückemeyer, G., Ruparel, R.: SecureBPMN: Modeling and enforcing access control requirements in business processes. In: *SACMAT*, pp. 123–126. ACM (2012)
- [18] Brucker, A.D., Petritsch, H.: Extending access control models with break-glass. In: Carminati, B., Joshi, J. (eds.) *SACMAT*, pp. 197–206. ACM (2009)
- [19] Casola, V., Fasolino, A., Mazzocca, N., Tramontana, P.: An ahp-based framework for quality and security evaluation. In: *CSE*, vol. 3, pp. 405–411 (2009)
- [20] Chan, H., Chieu, T., Kwok, T.: Autonomic ranking and selection of web services by using single value decomposition technique. In: *ICWS*, pp. 661–666 (2008)
- [21] Chan, S.W.: *Security annotations and authorization in glassfish and the Java EE 5 SDK* (2006)
- [22] Cherdantseva, Y.: *Secure*BPMN – a graphical extension for bpmn 2.0 based on a reference model of information assurance & security*. Ph.D. thesis, Cardiff University (2014)
- [23] Christensen, E., Curbera, F., Meredith, G., Weerawarana, S.: *Web services description language (WSDL) 1.1*. Tech. rep., W3C (2001)
- [24] Compagna, L., Guillemot, P., Brucker, A.D.: Business process compliance via security validation as a service. In: Oriol, M., Penix, J. (eds.) *Testing Tools Track of ICST*. IEEE Computer Society (2013)
- [25] Dijkman, R.M., Dumas, M., Ouyang, C.: Semantics and analysis of business process models in BPMN. *Information & Software Technology* **50**(12), 1281–1294 (2008)
- [26] Elshaafi, H., McGibney, J., Botvich, D.: Trustworthiness monitoring and prediction of composite services. In: *ISCC*, pp. 580–587 (2012)
- [27] Erl, T.: *Service-Oriented Architecture: Concepts, Technology, and Design*. Prentice Hall PTR (2005)
- [28] Felderer, M., Büchler, M., Johns, M., Brucker, A.D., Breu, R., Pretschner, A.: Security testing: A survey. *Advances in Computers* **101**, 1–51 (2016). doi: 10.1016/bs.adcom.2015.11.003
- [29] Fielding, R.T.: *REST: architectural styles and the design of network-based software architectures*. Phd dissertation, University of California, Irvine (2000)
- [30] Gromoff, A., Kazantsev, N., Ponfilenok, M., Stavenko, Y.: Newer approach to flexible business architecture of modern enterprise. In: *ICEIS*, pp. 326–332 (2013)
- [31] Hale, M., Gamble, R.: Risk propagation of security slas in the cloud. In: *IEEE GLOBECOM*, pp. 730–735 (2012)
- [32] Hale, M., Gamble, R.: Secagreement: Advancing security risk calculations in cloud services. In: *IEEE World Congress on Services*, pp. 133–140 (2012)
- [33] Henning, R.: Security service level agreements: Quantifiable security for the enterprise? In: *NSPW*, pp. 54–60 (2009)
- [34] Howard, M., Lipner, S.: *The Security Development Lifecycle*. Microsoft Press, Redmond, WA, USA (2006)

- [35] Irvine, C., Levin, T.: Quality of security service. In: NSPW, pp. 91–99 (2001)
- [36] Jaeger, M., Rojec-Goldmann, G., Muhl, G.: Qos aggregation in web service compositions. In: IEEE Int. Conf. on e-Technology e-Commerce and e-Service, p. 181185 (2005)
- [37] Jorstad, N., Landgrave, T.S.: Cryptographic algorithm metrics. In: Information Systems Security Conf. (1997)
- [38] Jürjens, J., Rumm, R.: Model-based security analysis of the german health card architecture. *Methods Inf Med* **47**(5), 409–416 (2008)
- [39] Krafzig, D., Banke, K., Slama, D.: Enterprise SOA: Service Oriented Architecture Best Practices. Prentice Hall (2005)
- [40] Lodderstedt, T., Basin, D.A., Doser, J.: SecureUML: a UML-based modeling language for model-driven security. In: Jézéquel, J.M., Hussmann, H., Cook, S. (eds.) UML, no. 2460 in LNCS, pp. 426–441. Springer (2002)
- [41] Marienfeld, F., Höfig, E., Bezzi, M., Flügge, M., Pattberg, J., Serme, G., Brucker, A.D., Robinson, P., Dawson, S., Theilmann, W.: Service levels, security, and trust. In: Barros, A., Oberle, D. (eds.) Handbook of Service Description: USDL and its Methods, chap. 12, pp. 295–326. Springer (2012)
- [42] Mauw, S., Oostdijk, M.: Foundations of attack trees. In: ICISC, pp. 186–198. Springer-Verlag, Berlin, Heidelberg (2005). doi: 10.1007/11734727_17
- [43] Modica, G.D., Petralia, G., Tomarchio, O.: A business ontology to enable semantic matchmaking in open cloud markets. In: SKG, pp. 96–103 (2012)
- [44] Mülle, J., von Stackelberg, S., Böhm, K.: A security language for BPMN process models. Tech. rep., KIT (2011)
- [45] OASIS: eXtensible Access Control Markup Language (XACML), version 2.0 (2005)
- [46] OMG: BPMN, version 2.0 (2011)
- [47] Organization for the Advancement of Structured Information Standards: Web services business process execution language (BPEL), version 2.0 (2007)
- [48] Paja, E., Dalpiaz, F., Poggianella, M., Roberti, P., Giorgini, P.: Modelling security requirements in socio-technical systems with STS-tool. In: Kirikova, M., Stirna, J. (eds.) CAiSE Forum, vol. 855, pp. 155–162 (2012)
- [49] Salnitri, M., Dalpiaz, F., Giorgini, P.: Designing secure business processes with SecBPMN. *Software & Systems Modeling* (2015)
- [50] Skoutas, D., Sacharidis, D., Simitsis, A., Kantere, V., Sellis, T.: Top-k dominant web services under multi-criteria matching. In: EDBT, pp. 898–909 (2009)
- [51] Sohr, K., Ahn, G.J., Gogolla, M., Migge, L.: Specification and validation of authorisation constraints using UML and OCL. In: di Vimercati, S.D.C., Syverson, P.F., Gollmann, D. (eds.) ESORICS, LNCS, vol. 3679, pp. 64–79. Springer (2005)
- [52] Vedamuthu, A.S., Orchard, D., Hirsch, F., Hondo, M., Yendluri, P., Boubez, T., Ümit Yalçınalp: Web services policy 1.5 (2007). URL <http://www.w3.org/TR/ws-policy/>
- [53] Wolter, C., Meinel, C.: An approach to capture authorisation requirements in business processes. *Requir. Eng.* **15**(4), 359–373 (2010)
- [54] Zhou, B., Arabo, A., Drew, O., Llewellyn-Jones, D., Merabti, M., Shi, Q., Waller, A., Craddock, R., Jones, G., Arnold, K.L.Y.: Data flow security analysis for system-of-systems in a public security incident. In: ACSF, pp. 8–14 (2008)
- [55] Zhou, B., Drew, O., Arabo, A., Llewellyn-Jones, D., Kifayat, K., Merabti, M., Shi, Q., Craddock, R., Waller, A., Jones, G.: System-of-systems boundary check in a public event scenario. In: SoSE (2010)
- [56] Zhou, B., Llewellyn-Jones, D., Shi, Q., Asim, M., Merabti, M., Lamb, D.: Secure service composition adaptation based on simulated annealing. In: ACSAC, pp. 49–55 (2012)