

Tisa: A Language Design and Modular Verification Technique for Temporal Policies in Web Services*

Hridesh Rajan¹, Jia Tao¹, Steve Shaner¹, and Gary T. Leavens²

¹ Iowa State University, Ames, Iowa, USA

{hridesh, jtao, smshaner}@iastate.edu

² University of Central Florida, Orlando, Florida, USA

leavens@eecs.ucf.edu

Abstract. Web services are distributed software components, that are decoupled from each other using interfaces with specified functional behaviors. However, such behavioral specifications are insufficient to demonstrate compliance with certain temporal non-functional policies. An example is demonstrating that a patient's health-related query sent to a health care service is answered only by a doctor (and not by a secretary). Demonstrating compliance with such policies is important for satisfying governmental privacy regulations. It is often necessary to expose the internals of the web service implementation for demonstrating such compliance, which may compromise modularity. In this work, we provide a language design that enables such demonstrations, while hiding majority of the service's source code. The key idea is to use greybox specifications to allow service providers to selectively hide and expose parts of their implementation. The overall problem of showing compliance is then reduced to two subproblems: whether the desired properties are satisfied by the service's greybox specification, and whether this greybox specification is satisfied by the service's implementation. We specify policies using LTL and solve the first problem by model checking. We solve the second problem by refinement techniques.

1 Introduction

Web services promote abstraction, loose coupling and interoperability of clients and services [1]. The key idea of web services is to introduce a published interface (often a description written in an XML-based language such as WSDL [2]), for communication between services and clients [1]. By allowing components to be decoupled using a specified interface, web services enable platform-independent integration.

Behavioral Contracts for Web Services. A behavioral contract for a web service specifies, for each of the web service's methods the relationships between its inputs and outputs. Such a contract treats the implementation of the service as a black box, hiding all the service's internal states from its clients. The benefit of this encapsulation is that clients do not depend upon the service's changeable design decisions. To illustrate, consider a healthcare service that allows patients to make appointments and ask prescription and health-related questions from healthcare practitioners [3].

* Rajan and Tao were supported in part by the NSF grant CNS 06-27354. Rajan, Shaner and Leavens were supported in part by the NSF grant CNS 08-08913.

An example JML-like contract [4] for such a service follows.

```

service Patient {
  /*@ requires pId >= 0; ensures result >=0; @*/
  int query(int pId, int msg);
  /*@ requires qId >= 0; ensures result >=0; @*/
  int retrieve(int qId);
}

```

The service description in this contract is written in a form similar to our language, Tisa, to make comparisons easier. It specifies that a service named `Patient` makes two web-methods available: `query` and `retrieve`. The `query` method takes a patient identifier and a message as arguments. The message is represented as an integer for simplicity (think of it as an index into a table of pre-defined questions, such as “does the test show I have AIDS?”). The precondition of calling this web-method is that the patient identifier is positive; the postcondition is that it returns a positive result. The `retrieve` method takes a query identifier as argument; its precondition is that this identifier must be positive. Its postcondition is that the result is also positive. These contracts could be checked by observing the interface of the web-methods [5,6,7,8,9].

Demonstrating Compliance to Temporal Policies. Let us now consider the following policy inspired from Barth *et al.*’s work [3]: “a health question about a patient should only be answered by the doctor”, “furthermore such answers should only be disclosed to the concerned patients”. We will refer to these as “HIPAA policies” as they are similar to regulations in the US health insurance portability and accountability act (HIPAA). The behavioral contract above is insufficient for demonstrating compliance with the HIPAA policies, as it does not provide sufficient details about the internal state of the service. For example, the entity that is finally receiving the query is hidden by `query`’s contract. Demonstrating compliance to such policies is important. In our example, a patient may feel much better about their queries regarding an AIDS test result, if such compliances were demonstrated by the service.

Compliance and Modularity at Conflict. Alternatively suppose the implementation of the two web-methods `query` and `retrieve` were available, including the component services that they use. Then demonstrating compliance to the two HIPAA policies would be equivalent to ensuring that the implementation avoids non-compliant states. However, by making code for these methods available, clients might write code that depends on implementation design decisions. As a result, changing these design decisions will become harder, as these changes could break client’s code [10].

We thus believe that, for web services, modularity [10] and verification of temporal policies are fundamentally in conflict. To make the service implementation evolvable, modularity requires hiding the design decisions that are likely to change. But to demonstrate compliance to key temporal policies, internal states need to be exposed.

A Language Design and Verification Logic. To reconcile these requirements, we propose a technique based on greybox specifications [11] that exposes only some internal states. This technique enables web service providers to demonstrate compliance to temporal policies, such that above, by exposing only parts of their implementation. A client can verify that the service complies with the desired policies by inspecting a greybox

```

1 service Secretary {
2   int query(int pId, int msg) {
3     preserve pId > 0 && msg > 0;
4     if (msg >= 2) {
5       query(pId,msg)@Doctor
6     }
7     else {
8       /* Appointment? */
9       establish result > 0
10    }
11  }
12  int retrieve(int qId) {
13    requires qId > 0 ensures result > 0
14  }
15 }

16 service Doctor {
17   int query(int pId, int msg) { /* Re: Test */
18     requires pId > 0 && msg >= 2 ensures result > 0
19   }
20   int retrieve(int qId) {
21     requires qId > 0 ensures result > 0
22   }
23 }
24 service Patient {
25   int query(int pId, int msg) {
26     query(pId, msg)@Secretary;
27   }
28   int retrieve(int qId) {
29     preserve qId > 0;
30     if ((qId/1000)==1) { retrieve(qId)@Secretary}
31     else if ((qId/1000)==2) { retrieve(qId)@Doctor}
32   }

```

Fig. 1. An Example Greybox Specification

specification. Providers can also choose to hide many implementation details, so the service’s implementation can evolve as long as it refines the specification [12,13].

To illustrate, consider the greybox specification shown in Figure 1. This example has three services. In each service the methods are web-methods that may be called by clients and other services. Specification expressions of the form **preserve** e , **establish** e , and **requires** e_1 **ensures** e_2 are used within these methods to hide internal details. The code that is not hidden by specification expressions is exposed. Calls to web-methods are written using an at-sign (@), such as `query(pId, msg)@Secretary`. For simplicity, Tisa only allows integers to be passed as arguments in such remote calls, thus we encode questions using integers: 1 for appointments, 2 for prescriptions, and higher numbers for health-related questions. Contrary to standard black box specifications, internal states of the service, including calls to other services are exposed. By analyzing lines 26 and 4–6 (in that order) one could conclude that “health questions by patients are answered by the doctor.” Demonstrating compliance to temporal policies thus becomes possible. Note that this specification only exposes selected details about the implementation. For example, the specification of `retrieve` on line 13 hides all details of how this service responds to appointment questions. Therefore, it hides the design decisions made in the implementation of creating, storing, and forwarding responses.

Contributions. An important contribution is the identification of the conflict between verification of temporal policies and modularity in web services. We show how to resolve this conflict using greybox specifications. Our language, *Tisa*, supports specification of policies specified in a variant of linear temporal logic [14], greybox specification [11] and a simple notion of refinement [12,13,15] for modular reasoning about correctness of implementations with respect to such policies. As usual, implementations are hidden, but policies and greybox specifications are public. To demonstrate these claims, we present two preliminary verification techniques: one checks if a greybox specification satisfies a temporal policy, the second checks whether a service implementation refines its greybox specification. (The first technique could be used by the clients to select a service whose specification satisfies their desired policies.) We also show soundness:

```

program ::= decl* client
decl ::= classdecl | servicedecl
classdecl ::= class c extends d { field* meth* }
servicedecl ::= service w { field* meth* }
client ::= client w { e }
field ::= t f ;
meth ::= t m (form*) { e }
form ::= t var, where var ≠ this and var ≠ thisSite
t ::= c | int
e ::= n | e == e | e != e | e > e | e < e | e >= e | e <= e
      | e + e | e - e | e * e | ! e | e && e | e ' | ' e | isNull(e)
      | if (e) { e } else { e } | new c () | var
      | null | e.m(e*) | e.f | e.f = e | cast c e | form = e ; e
      | e ; e | w | m(e*) @e | refining spec { e }

```

$n \in \mathcal{N}$, the set of numeric, integer literals
 $c, d \in \{\text{Object}, \text{Site}\} \cup \mathcal{C}$,
 \mathcal{C} is the set of class names
 $f \in \mathcal{F}$, the set of field names
 $m \in \mathcal{M}$, the set of method names
 $\text{var} \in \{\text{this}, \text{thisSite}\} \cup \mathcal{V}$,
 \mathcal{V} is the set of variable names
 $w \in \mathcal{W} \subseteq \mathcal{C}$,
 \mathcal{W} is the set of web service names

Fig. 2. Abstract syntax, based on [25, Figure 3.1, 3.7]

that the composition of these two verification techniques, applied modularly by clients and all service providers, implies that the web service implementation satisfies the specified temporal policies. In practice, some additional technique, such as proof-carrying code [16], or a hardware-based root of trust [17,18] would be needed to satisfy clients that web services in fact satisfy their specifications.

2 Tisa Language Design

In this section, we describe Tisa, an object-oriented (OO) language that incorporates ideas from existing work on specification languages, web services authentication languages and modeling languages. In particular, Tisa’s design is inspired by Argus [19] and the work of Gordon and Pucella [20]. (Furthermore, some of our descriptions of the language syntax are adapted from Ptolemy [21].) Tisa is a distributed programming language with statically created web services and a single client, each of which has its own address space. Web services are named and declare web-methods, which can be called by the client and by other services. As a small, core language, the technical presentation of Tisa shares much in common with MiniMAO₁ [22], a variant of Featherweight Java [23] and Classic Java [24]. Tisa has classes, objects, inheritance, and subtyping, but it does not have **super**, interfaces, exception handling, built-in value types, privacy modifiers, or abstract methods. Furthermore, other features of web-service description languages (WSDLs) such as composite data types for exchanging messages between services, messages, ports, one-way vs. request-response operations, etc. are omitted to avoid complications in Tisa’s theory. However, most of these are syntactic sugars that can be desugared to existing constructs in Tisa. Tisa features new mechanisms for declaring policies and greybox specifications. Our description starts with its programming features, and then describes its specification features.

2.1 Program Syntax

The syntax of Tisa executable programs is shown in Figure 2 and explained below. A Tisa program consists of zero or more declarations, and a client (see Figure 3). Declarations are either class declarations or web service declarations.

```

1 class Query extends Object {
2   int pId; int msg; int qId;
3 }
4 class Queue extends Object { //...
5   int add(int pId, int msg, int qId){
6     /* add to inner list */; qId
7   }
8   service Secretary {
9     Queue queryQ; Hashtable responses;
10    int ticket; Log log;
11    int query(int pId, int msg) {
12      refining preserve pId > 0 && msg > 0 {
13        log.recordCurrentTime()
14      };
15      if (msg >= 2) {
16        query(pId, msg)@Doctor
17      } else { /* Re: Appointment */
18        refining establish result > 0 {
19          ticket = ticket + 1;
20          queryQ.add(pId, msg, ticket + 1000)
21        } }
22    int respond(int qId,int pId,int msg){
23      /* Encode patient's information */
24      responses.add(qId, pId*1000 + msg);
25      queryQ.remove(qId)
26    }
27    int retrieve(int qId) {
28      refining requires qId > 0
29      ensures result > 0 {
30        responses.get(qId)
31      } }
32  }
33  service Doctor {
34    Queue topQ; Queue medQ; Queue lowQ;
35    int query(int pId, int msg) {
36      refining requires pId > 0 && msg >= 2
37      ensures result > 0 {
38        ticket = ticket + 1;
39        if (msg > 500) {
40          topQ.add(pId, msg, ticket + 2000)
41        } else if (msg > 250) {
42          medQ.add(pId, msg, ticket + 2000)
43        } else {
44          lowQ.add(pId, msg, ticket + 2000)
45        };
46      }
47      /* retrieve similar to Secretary's */
48    }
49    service Patient {
50      int query(int pId, int msg) {
51        query(pId, msg)@Secretary
52      }
53      int retrieve(int qId) {
54        if ((qId/1000) == 1) {
55          retrieve(qId)@Secretary
56        } else if ((qId/1000) == 2) {
57          retrieve(qId)@Doctor
58        } }
59      client User{
60        int qid = query(101,3)@Patient;
61        retrieve(qid)@Patient
62      }

```

Fig. 3. An Example Tisa Implementation

Each web service has a name (w) representing that web service; thus web service names can be thought of as web sites. (The mapping of web services to actual computers is not specified in the language itself.) A web service can be thought of as a singleton object; however, each web service has a separate address space and its methods can only be called using a remote procedure call.

An example web service declaration for the service `Patient` appears on lines 49–62 in Figure 3. This service contains two web-methods declaration, named `query` and `retrieve`. The web-method `query` takes a patient Id and message as arguments and returns a unique query Id generated according to the input arguments. The web-method `retrieve` takes query Id as an argument and returns an answer message which encodes a patient Id. A client declares a name and runs an expression that is the main expression of the program. We next explain class declarations and expressions.

Class Declarations. Class declarations may not be nested. Each class has a name (c) and names its superclass (d), and may declare finite number of fields ($field^*$) and methods ($meth^*$). Field declarations are written with a class name, giving the field's type, followed by a field name. Methods also have a C++ or Java-like syntax, although their body is an expression.

Expressions. Tisa is an expression language. Thus the syntax for expressions includes integer literals, various standard integer and logical operations, several standard OO expressions and also some expressions that are specific to web services. The logical

<pre> specification ::= servicespec* servicespec ::= service w { wmspec* } wmspec ::= t m (form*) { se } form ::= t var, where var ≠ thisSite spec ::= requires sp ensures sp </pre>	<pre> se ::= sp spec se; se form = se; se m (sp*) @sp if (sp) { se } else { se } sp ::= n sp == sp sp != sp sp > sp sp < sp sp >= sp sp <= sp sp + sp sp - sp sp * sp ! sp sp && sp sp ' ' sp var w </pre>
---	--

Fig. 4. Syntax for Writing Specifications in Tisa

operations operate on integers, with 0 representing false, and all other integer values representing true. An **if** (e_1) { e_2 } **else** { e_3 } expression tests if e_1 is non-zero; if so it returns the value of e_2 , otherwise it returns the value of e_3 .

The standard OO expressions include object construction (**new** $c()$), variable dereference (var , including **this**), field dereference ($e.f$), **null**, cast (**cast** $t e$), assignment to a field ($e_1.f = e_2$), sequencing ($e_1; e_2$), casts and a definition block ($t var = e_1; e_2$). The other OO expressions are standard [25,22].

There are three new expressions: web service names, web-method calls, and refining statements. Web service names of form w are constants. A *web-method call* has the form ($m(e^*) @e_w$), where the expression following the at-sign (e_w) denotes the name of the web service name that will execute the web-method call named m with formals e^* . A **refining** statement, of the form **refining** $spec$ { e }, is used in implementing Tisa's greybox specifications (see below). It executes the expression e , which is supposed to satisfy the specification $spec$.

2.2 Specification Constructs

The syntax for writing specifications in Tisa is shown in Figure 4. In this figure, all nonterminals that are used but not defined are the same as in Figure 2. Specifications consist of several service specifications (*servicespec*). (Since we only permit integers to be sent to and returned from web-method calls, we omit class declarations from specifications.) A service specification may contain finite number of web-method specifications (*wmspec*). All fields are hidden, so field declarations are not allowed in a service specification. The body of a web-method specification contains a side-effect free expression (se). Many expressions from Figure 2 also appear as such side-effect free expressions, but not field-related operations, method calls, and **isNull**. Web-method call expressions are allowed and so are local variable definition expressions.

The main new feature of specifications, borrowed from the refinement calculus and the greybox approach, is the specification expression ($spec$). Such an expression hides (abstracts from) a piece of code in a correct implementation. The most general form of specification expression is **requires** sp_1 **ensures** sp_2 , where sp_1 is a precondition expression and sp_2 is a postcondition. Such a specification expression hides program details by specifying that a correct implementation contains a **refining** expression whose body expression, when started in a state that satisfies sp_1 , will terminate in a state that satisfies sp_2 [15].

In examples we use two sugared forms of specification expression. The expression **preserve** sp is sugar for **requires** sp **ensures** sp and **establish** sp is sugar for **requires** 1 **ensures** sp .

An example greybox specification of the web service `Patient` appears in Figure 1. The specification of the web-method `query` appears on line 26, and specifies (and

thus exposes) all the code for that method. The specification of `retrieve` hides a bit more in its **preserve** expression (line 29). But it also exposes code that makes a web-method call `retrieve` to the `Secretary` or `Doctor`. With these greybox specifications, enough details are exposed about what the service does when invoking other services, which makes it feasible to show compliance to the HIPAA policies.

2.3 Constructs for Specifying Policies

Our simple policy specification language is similar to Linear Temporal Logic [14].

$$\mathcal{P}(\text{specification}) ::= \mathcal{P}(\text{specification}) \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \mathbf{U} \phi_2 \mid \mathbf{X} \phi$$

The language specifies histories that are sequences of web method calls. For a given *specification*, a policy can be an atomic proposition in $\mathcal{P}(\text{specification})$; a negation of a policy or boolean combination of policies. For simplicity here we take the set of legal propositions $\mathcal{P}(\text{specification})$ to be all legal web-method calls in the given specification. This set can be statically computed from the specification against which the policy is to be verified by traversing the abstract syntax tree of the specification up to the depth of web-method specifications. The operator **U** is read as “until” and **X** as “next.” $\phi_1 \mathbf{U} \phi_2$ states that policy ϕ_2 must be satisfied after policy ϕ_1 is satisfied along all executions of the service. $\mathbf{X}\phi$ states that policy ϕ must be satisfied in the next state (i.e., at the next web method call). We also use the following common abbreviations:

$$\begin{array}{lll} \phi_1 \vee \phi_2 \equiv \neg(\neg\phi_1 \wedge \neg\phi_2) & \phi_1 \rightarrow \phi_2 \equiv \neg\phi_1 \vee \phi_2 & \text{true} \equiv \phi \vee \neg\phi \\ \text{false} \equiv \neg\text{true} & \mathbf{F}\phi \equiv \text{true} \mathbf{U} \phi & \mathbf{G}\phi \equiv \neg\mathbf{F}\neg\phi \end{array}$$

The constant *true* means that the service does not have any obligation. The operator **F** is read as “eventually” and **G** as “always”. Below we present two sample policies for our healthcare service example.

$$\begin{aligned} \phi_1 &= \mathbf{G}(\text{query@Patient} \wedge (\mathbf{X}\mathbf{F}(\text{query@Secretary} \vee \mathbf{X}\mathbf{F}\text{query@Doctor}))) \\ \phi_2 &= \mathbf{G}(\text{retrieve@Patient} \wedge \mathbf{X}\mathbf{F}\text{retrieve@Doctor} \rightarrow \neg \mathbf{X}\mathbf{F}\text{retrieve@Secretary}) \end{aligned}$$

The policy ϕ_1 states that whenever there is a web-method call `query@Patient`, there is eventually a web-method call `query` at one of the sites `Secretary` or `Doctor`. This policy says that a query is eventually delivered to one of the healthcare providers. The policy ϕ_2 encodes the constraint that a health answer that comes from doctors goes directly to the patient, and is never forwarded to secretaries. In terms of the service specification, if there is a web-method call `retrieve@Patient` and it is followed by a web-method call `retrieve@Doctor`, then there is never a web-method call `retrieve` at the site `Secretary` in the same trace.

2.4 Dynamic Semantics of Tisa’s Constructs

This section defines a small step operational semantics for Tisa programs (adapted from Clifton’s work [25]). In the semantics, all declarations are formed into a single class table that maps class names and web service names to class and service declarations, respectively. However, despite this global view of declarations, the model of storage is distributed, with each web service having an independent store.

Evaluation relation: $\hookrightarrow: \Gamma \rightarrow \Gamma$

(WEB METHOD CALL)

$$\frac{\begin{array}{l} \Pi = \{var_i : \mathbf{var} \ t_i \mid 1 \leq i \leq n\} \uplus \{\mathbf{this} : \mathbf{var} \ c_2\} \uplus \{\mathbf{thisSite} : \mathbf{var} \ Site\} \quad \nu = \mathbf{frame} \ \rho \ \Pi \\ \rho = \{var_i \mapsto v_i \mid 1 \leq i \leq n\} \oplus (\mathbf{this} \mapsto loc) \oplus (\mathbf{thisSite} \mapsto w) \\ (loc, c_2, t \ m(t_1 var_1, \dots, t_n var_n)\{e\}) = \mathit{find}(w, m) \end{array}}{\langle \mathbb{E}[m(v_1, \dots, v_n) @ w], J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under} \ e], \nu + J, S \rangle}$$

(REFINING)

$$\frac{n \neq 0}{\langle \mathbb{E}[\mathbf{refining} \ \mathbf{requires} \ n \ \mathbf{ensures} \ e' \ (e'')], J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{evalbody} \ e'' \ e'], J, S \rangle}$$

(EVALBODY)

$$\frac{\begin{array}{l} \rho = \mathit{envOf}(\nu) \quad \Pi = \mathit{tenvOf}(\nu) \quad w = \mathit{thisSite}(\nu) \quad t = \mathit{typeOf}(v, S, w) \\ \rho' = \Pi \uplus \{\mathbf{result} : v\} \quad \Pi' = \Pi \uplus \{\mathbf{result} : \mathbf{var} \ t\} \quad \nu' = \mathbf{frame} \ \rho' \ \Pi' \end{array}}{\langle \mathbb{E}[\mathbf{evalbody} \ v \ e'], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under} \ \mathbf{evalpost} \ v \ e'], \nu' + \nu + J, S \rangle}$$

(EVALPOST)

$$\frac{n \neq 0}{\langle \mathbb{E}[\mathbf{evalpost} \ v \ n], J, S \rangle \hookrightarrow \langle \mathbb{E}[v], J, S \rangle}$$

(UNDER)

$$\langle \mathbb{E}[\mathbf{under} \ v], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[v], J, S \rangle$$

Fig. 5. Operational semantics of Tisa. Standard OO rules are omitted.

The operational semantics relies on four expressions, not part of Tisa's surface syntax, to record final or intermediate states of the computation. The *loc* expression represents locations in the store. The **under** expression is used as a way to mark when the evaluation stack needs popping. The **evalbody** and **evalpost** are used in evaluation of specification expressions. The three exceptions `NullPointerException`, `ClassCastException`, and `SpecException` record various problems orthogonal to the type system.

A configuration in the semantics contains an expression (e), an evaluation stack (J), and a store (S). The current web service name is maintained in the evaluation stack under the name **thisSite**. The auxiliary function *thisSite* extracts the current web service name from a stack frame. Stacks are an ordered list of frames, each frame recording the static environment, ρ , and a type environment. (The type environment, Π , is only used in the type soundness proof.) The static environment ρ maps identifiers to values. A value is a number, a web service name (site), a location, or **null**. Stores are maps from locations to storable values, which are object records. Object records have a class and also a map from field names to values.

The semantics is presented as a set of evaluation contexts \mathbb{E} and an one-step reduction relation [26] that acts on the position in the overall expression identified by the evaluation context as shown in Figure 5. Standard OO rules are presented in our technical report [27]. The key rule is (WEB METHOD CALL), which uses the auxiliary function *find* to retrieve the body of the web method from a class table CT implicitly used by the semantics. It creates the frame for execution of the web method with necessary static environment and type environment and starts execution of the web method body. The **under** e expression is used in the resulting configuration to mark that the stack should be popped when the evaluation of e is finished.

Evaluation of a **refining** expression involves 3 steps. First the precondition is evaluated (due to the context rules). If the precondition is non-zero (i.e., true), then the next configuration is **evalbody** $e'' \ e'$, where e'' is the body and e' is the postcondition

(regarded as an expression). The body is then evaluated; if it yields a value v , then the next configuration is **under** `evalpost` v e' , with a new stack frame that binds **result** to v pushed on the stack. The type of **result** in the type environment Π' is determined by the auxiliary function *typeOf*. Finally, the (EVALPOST) rule checks that the postcondition is true and uses the body's value as the value of the expression.

3 Verification of Policies in Tisa

A key contribution of our work is to decouple, with Tisa's language design, the verification of whether a policy is satisfied by a web service implementation into two verification tasks that can proceed modularly and independently. The first task is to verify whether a policy is satisfied by the service specification. The second task is to verify whether the service specification is satisfied by the service implementation. Three benefits follow from this modular approach. First, the service implementation need not be visible to clients, as a client uses the specification to determine whether their desired policies hold. Thus, our approach achieves modularity for service implementations. Second, regardless of the number of clients, the second verification task must only be done once; thus our approach is likely to be scalable for web service providers. Last but not the least, policy verification is performed on the (generally smaller) specification. Thus, our approach has efficiency benefits for policy verification.

Determining whether a policy is satisfied by the specification can be reduced to a standard model checking problem [14]. We claim no contribution here; rather, the novelty of our approach is in a combination of these two techniques, enabled by a careful language design. To show the feasibility of applying ideas from model checking [14] and refinement calculus [12,13] to our problem, in the rest of this section we describe our techniques for verifying policies and refinement.

3.1 Verifying Policies

We adopt the standard automata-theoretic approach for verifying linear temporal logic formulas proposed by Vardi and Wolper [28] to verify policies in Tisa. Following Vardi and Wolper [28], a policy $\phi \in \Phi(\mathcal{S})$ is viewed as a finite-state acceptor and a specification \mathcal{S} as a finite-state generator of expression execution histories. Thus the specification \mathcal{S} satisfies policy ϕ if every (potentially infinite) history generated by \mathcal{S} is accepted by ϕ , in other words, if $\mathcal{S} \cap \neg\phi$ is empty.

Figure 6 shows main parts of an algorithm for constructing a finite-state machine $\mathcal{F}(\mathcal{S}) = (\mathcal{Z}, z_0, R, \Delta)$ from a Tisa specification \mathcal{S} . Here, \mathcal{Z} is a finite set of states, z_0 is the initial state, R is a total accessibility relation, $\Delta : \mathcal{Z} \rightarrow 2^{\mathcal{P}(\mathcal{S})}$, which determines how truth values are assigned to propositions in each state [28, pp. 5]. All rules make use of unions for joining set of states (\mathcal{Z}) and disjoint union (\uplus) for joining propositions. Rules for standard OO expressions are omitted.

The (IF EXP FSM) rule demonstrates creation of non-deterministic transitions in the state machine. It computes the FSMs corresponding to the true branch and the false branch of the **if** expression with initial states z' and z'' and joins these two FSMs to make a new FSM with initial state z . Corresponding to the state z' , which corresponds

Production relation: $NT \vdash se \rightsquigarrow (Z, z_0, R, \Delta), NT$ where $NT \in \mathcal{NT} = \mathcal{W} \times \mathcal{M} \rightarrow \mathcal{Z}$

(IF EXP FSM)

$$\frac{NT \vdash se' \rightsquigarrow (Z', z', R', \Delta'), NT' \quad NT' \vdash se'' \rightsquigarrow (Z'', z'', R'', \Delta''), NT'' \quad Z = Z' \cup Z'' \cup \{z\} \quad \Delta = \Delta' \uplus \Delta'' \uplus \{(z', \{sp\}), (z'', \{!sp\})\} \quad R = R' \cup R'' \cup \{(z, z'), (z, z'')\}}{NT \vdash \mathbf{if} (sp) \{se'\} \mathbf{else} \{se''\} \rightsquigarrow (Z, z, R, \Delta), NT''}$$

(WEB METHOD CALL FSM 1)

$$\frac{NT' = NT \cup ((w, m), z) \quad m(t_1, \dots, t_n)\{se\} = \mathit{find}(w, m) \quad NT' \vdash se \rightsquigarrow (Z', z', R', \Delta'), NT'' \quad Z = Z' \cup \{z\} \quad \Delta = \Delta' \uplus \{(z', \{m@w\})\} \quad R = R' \cup \{(z, z')\}}{NT \vdash m(v_1, \dots, v_n)@w \rightsquigarrow (Z, z, R, \Delta), NT''}$$

(WEB METHOD CALL FSM 2)

$$\frac{z = NT(w, m)}{NT \vdash m(v_1, \dots, v_n)@w \rightsquigarrow (\{z\}, z, \{\}, \{\}), NT}$$

(SPEC EXP FSM)

$$\frac{Z = \{z_1, z_2, z_3, z_4\} \quad R = \{(z, z_1), (z, z_2), (z_1, z_3), (z_1, z_4), (z_3, z')\} \quad \Delta_{pre} = \{(z_1, \{sp_1\}), (z_2, \{!sp_1\})\} \quad \Delta = \Delta_{pre} \uplus \{(z_3, \{sp_1, sp_2\}), (z_4, \{sp_1, !sp_2\})\}}{NT \vdash \mathbf{requires} \ sp_1 \ \mathbf{ensures} \ sp_2 \rightsquigarrow (Z, z, R, \Delta), NT}$$

Fig. 6. Finite-state machine construction, built from expressions in a specification

to the true branch, the proposition sp is added to Δ , which corresponds to the conditional expression evaluating to the truth value true. Similarly for the state z'' , which corresponds to the false branch, the proposition $!sp$ is added to Δ , which corresponds to the conditional expression evaluating to the truth value false.

The (SPEC EXP FSM) rule models the cases for satisfaction of precondition and postcondition. The (WEB METHOD CALL FSM) rules make use of a table NT that maps pairs of web service names and method names (w, m) to states. This table is used to account for recursion in web-method calls. Finally, the finite-state machine for a service specification is created by first creating finite-state machines for each of its web-method specifications as if it is being called and by joining them using an extra state that becomes the new initial state.

Given the FSM $\mathcal{F}(\mathcal{S})$ we construct a Büchi automaton [29], $\mathcal{B}(\neg\phi)$ for the policy $\phi \in \Phi(\mathcal{S})$ as shown by Vardi and Wolper [28]. Specification \mathcal{S} satisfies the policy ϕ if $\mathcal{F}(\mathcal{S}) \cap \mathcal{B}(\neg\phi)$ is empty.

3.2 Verifying Refinement

Our technique for checking whether a program refines a specification in Tisa is similar to the work of Shaner, Leavens and Naumann [15]. An implementation refines a specification if it meets two criteria: first, that the code and specification are structurally similar and second, that the body of every **refining** expression obeys the specification it is refining. By structural similarity we mean that for every non-specification expression in the specification, the implementation has the identical expression at that position in the code. This is checked in a top-down manner as shown in Figure 7. The operational semantics rules (REFINING), (EVALBODY) and (EVALPOST) ensure that the body of every **refining** expression obeys the specification it is refining.

$$\begin{array}{c}
\text{(PROGRAM REF)} \\
\frac{\forall i \in \{1..m\} \exists j \in \{1..n\} \text{decl}_j \in \text{servicedecl} \wedge \text{servicespec}_i \sqsubseteq \text{decl}_j}{\text{servicespec}_1 \dots \text{servicespec}_m \sqsubseteq \text{decl}_1 \dots \text{decl}_n}
\end{array}
\qquad
\begin{array}{c}
\text{(SP REF)} \\
\frac{sp = e}{sp \sqsubseteq e}
\end{array}$$

$$\begin{array}{c}
\text{(SERVICE REF)} \\
\frac{\forall i \in \{1..m\} \exists j \in \{1..n\} \text{wmspec}_i \sqsubseteq \text{meth}_j}{\text{service } w \{ \text{wmspec}_1 \dots \text{wmspec}_n \} \\ \sqsubseteq \text{service } w \{ \text{field}_1 \dots \text{field}_f \text{meth}_1 \dots \text{meth}_n \}}
\end{array}
\qquad
\begin{array}{c}
\text{(WEB METHOD REF)} \\
\frac{se \sqsubseteq e}{t \text{m}(\text{form}_1 \dots \text{form}_n) \{se\} \\ \sqsubseteq t \text{m}(\text{form}_1 \dots \text{form}_n) \{e\}}
\end{array}$$

$$\begin{array}{c}
\text{(SEQ EXP REF)} \\
\frac{se_1 \sqsubseteq e_1 \quad se_2 \sqsubseteq e_2}{se_1 ; se_2 \\ \sqsubseteq e_1 ; e_2}
\end{array}
\qquad
\begin{array}{c}
\text{(IF EXP REF)} \\
\frac{sp \sqsubseteq e_b \quad se_T \sqsubseteq e_T \quad se_F \sqsubseteq e_F}{\text{if } (sp) \{se_T\} \text{ else } \{se_F\} \\ \sqsubseteq \text{if } (e_b) \{e_T\} \text{ else } \{e_F\}}
\end{array}
\qquad
\begin{array}{c}
\text{(DEF EXP REF)} \\
\frac{sp \sqsubseteq e_{init} \quad se \sqsubseteq e_{body}}{form = sp ; se \\ \sqsubseteq form = e_{init} ; e_{body}}
\end{array}$$

$$\begin{array}{c}
\text{(WEBCALL EXP REF)} \\
\frac{(\forall i \in \{1..n\} :: sp_i \sqsubseteq e_i) \quad sp_w \sqsubseteq e_w}{m(sp_1, \dots, sp_n) @ sp_w \sqsubseteq m(e_1, \dots, e_n) @ e_w}
\end{array}
\qquad
\begin{array}{c}
\text{(SPEC EXP REF)} \\
\frac{(\text{requires } sp_1 \text{ ensures } sp_2) = spec}{\text{requires } sp_1 \text{ ensures } sp_2 \sqsubseteq \text{refining } spec \{e\}}
\end{array}$$

Fig. 7. Inference rules for proving Tisa refinement

3.3 Soundness of Verification Technique

The proof of soundness of our verification technique uses the following three definitions.

Definition 1 (A Path for S). Let S be a specification and $\mathcal{F}(S) = (Z, z_0, R, \Delta)$ be the FSM for S constructed using algorithm shown in Figure 6. A path t for S is a (possibly infinite) sequence of pairs $(z_i, \Delta(z_i))$ starting with pair $(z_0, \Delta(z_0))$, where for each $i \geq 0$, $z_i \in Z$ and $(z_i, z_{i+1}) \in R$.

Definition 2 (A Path for P). Let P be a program and $\mathcal{CFG}(P) = (Z', z'_0, R', \Delta')$ be an annotated control flow graph for P , where Z' is the set of nodes representing expressions in program, R' is the control flow relation between nodes, and $\Delta' : Z' \rightarrow 2^{\mathcal{P}(P)}$ is such that for each $z'_i \in Z'$, if it represents a web-method call expression $m(\dots)@w$ then $(z'_i, \{m@w\}) \in \Delta'$. A path t' for P is a (possibly infinite) sequence of pairs $(z'_i, \Delta(z'_i))$ starting with pair $(z'_0, \Delta(z'_0))$, where for each $i \geq 0$, $z'_i \in Z'$ and $(z'_i, z'_{i+1}) \in R'$.

Definition 3 (Path Refinement). Let t be a path for S and t' be a path for P . Then t is refined by t' , written $t \sqsubseteq t'$, just when one of the following holds:

- $t \equiv t'$ i.e., for each $i \geq 0$, $(z_i, \delta_i) \in t$ and $(z'_i, \delta'_i) \in t'$ implies $z_i = z'_i$ and $\delta_i = \delta'_i$,
- $t = (z, \delta) + t_1$ and $t' = (z', \delta') + t'_1$ and $\delta \Rightarrow \delta'$ and $t_1 \sqsubseteq t'_1$,
- $t = (z, \delta) + t_1$ and $t' = (z'_1, \delta'_1) + \dots + (z'_n, \delta'_n) + t'_1$ and $\delta \Rightarrow (\delta'_1 \uplus \dots \uplus \delta'_n)$ and $t_1 \sqsubseteq t'_1$, or
- $t = t_1 + t_2$ and $t' = t'_1 + t'_2$ and $t_1 \sqsubseteq t'_1$ and $t_2 \sqsubseteq t'_2$.

Lemma 1. Let $P \in \text{program}$ and $S \in \text{specification}$ be given. If P refines S , then for each path t' for P there exists a path t for S such that $t \sqsubseteq t'$.

Proof Sketch: The proof for this lemma follows from structural induction on the refinement rules shown in Figure 7. Details are contained in our technical report [27].

Lemma 2. *Given a specification \mathcal{S} and a policy $\phi \in \Phi(\mathcal{S})$, the automaton $\mathcal{F}(\mathcal{S}) \cap \mathcal{B}(\neg\phi)$ accepts a language, which is empty when the specification satisfies the policy.*

The proof of this lemma follows from standard proofs in model checking, in particular, from Lemma 3.1, Theorem 2.1 and Theorem 3.3. given by Vardi and Wolper [28, pp. 4,6]. Details are contained in our technical report [27].

Theorem 1. *Let \mathcal{S} be a specification, ϕ be a policy in $\Phi(\mathcal{S})$, and P be a program. Let ϕ be satisfied by the specification \mathcal{S} and P be a refinement of \mathcal{S} (as defined in Figure 7). Then the policy ϕ is satisfied by the program P .*

Proof Sketch: The proof follows from lemma 1 and 2. From lemma 1, we have that each path in the program refines a path in the specification. From lemma 2 and the assumptions of this theorem, we have that ϕ is satisfied on all paths in \mathcal{S} . Thus, ϕ , which is written over $\mathcal{P}(\mathcal{S})$, is also satisfied for P .

4 Related Work

In this section, we discuss techniques that are closely related to our approach.

Greybox specifications. We are not the first to consider greybox specifications [11] as a solution for verification problems. Barnett and Schulte [30,31] have considered using greybox specifications written in AsmL [32] for verifying contracts for .NET framework. Wasserman and Blum [33] also use a restricted form of greybox specifications for verification. Tyler and Soundarajan [34] and most recently Shaner, Leavens, and Naumann [15] have used greybox specifications for verification of methods that make mandatory calls to other dynamically-dispatched methods. Compared to these related ideas, to the best of our knowledge our work is the first to consider greybox specifications as a mechanism to decouple verification of web services without exposing all of their implementation details. Secondly, most of these, e.g. Shaner, Leavens, and Naumann [15] use the refinement of Hoare logic as their underlying foundation. This was insufficient to tackle the problem that we address, which required showing refinement of (a variant of) linear temporal logic. Thus adaptation of much of their work was not possible, although we were able to adapt the notion of structural refinement.

Specification and Verification Techniques for Web Services. The technique proposed by Bravetti and Zavattaro [35] for determining whether the behavioral contract of a service correctly refines its desired requirements in a composition of web-services is closely related and complementary to this work. The main difference between this work and the current work is that we verify refinement of greybox specifications by service implementations that allows us to reason about temporal policies, while hiding much of the implementation. However, we foresee a combination of our work and Bravetti and Zavattaro's work for determining fitness of a service implementation in a desired composition of web-services.

Some approaches have recently been proposed to verify contracts for web services, as seen in the works of Acciai and Boreale [36], Kuo *et al.* [8], Baresi *et al.* [6], Barbon *et al.* [5], etc. These ideas focus on verifying the behavioral contracts as defined by the externally visible interface of the web services, whereas our work provides a

principled, modular technique for verifying such policies that require inspecting the web service implementation to a limited extent.

Castagna, Gesbert and Padovani present a formalism for specifying web services based on the notion of “filtering” the possible behaviors of an existing web service to conform to the behavior of some contract [7]. These filters take the form of coercions that limit when and how an available service may be consumed. These coercions permits contract subtyping and support reasoning in a language-independent way about the sequence of reads and writes performed between service clients and providers. Their contracts are intended to constrain the usage scenarios of a web service, whereas the present work describes a modular way to specify the observable behaviors that occur inside service implementations.

Bartoletti *et al.* [37] provide a formalization of web service composition in order to reason about the security properties provided by connected services. While they ignore policy language details, our work shows how the amount of overhead used to relate specifications to policies depends on the level of detail in the policy language. Furthermore, we believe greybox reasoning grants real benefits in readability and modularity over their type system. We view later work developing executable specifications for design of web services [38] as possible future work for Tisa.

Another approach [39] proposes an architecture to enforce these access policies at component web services, but again the work is tightly coupled to the WS-SensFlow and Axis implementations. Srivatsa *et al.* [40] propose an Access Control system for composite services which does not take care of the Trust in the resulting service oriented architecture. Skalka and Wang [41] introduced a trust but verify framework which is an access control system for web services, but they do not provide temporal reasoning for the verification of policies. By recording the sequence of program events in temporal order, Skalka and Smith [42] are able to verify the policies such as whether the events were happened in a reasonable order, but the mechanism does not support decoupling the model and the implementation. Other approaches [43,44] either do not have a formal model supporting them or are tightly coupled with implementations.

Future Work and Conclusions

We have designed Tisa to be a small core language to clearly communicate how it allows users to balance compliance and modularity in web service specification. However, our desire for simplicity and clarity led us to leave for future work many practical and useful extensions. The most important future work in the area of Tisa’s semantics is to investigate refinement of information flow properties. It would also be interesting to investigate the utility of Tisa’s specification forms for reasoning about the composition of web services.

Verifying web services is an important problem [7,5,6,8,9], which is crucial for wider adoption of this improved modularization technique that enables new integration possibilities. There are several techniques for verifying web-services using behavioral interfaces, but none facilitates verification that requires access to internal states of the service. To that end, the key contribution of this work is to identify the conflict between verification of temporal properties and modularity requirements in web services. Our language design, Tisa, addresses these challenges. It allows service providers to demonstrate

compliance to policies expressed in an LTL-like language [14]. We also showed that policies in Tisa can be verified by clients using just the specification. Furthermore, refinement of specifications by program ensures that conclusion drawn from verifying policies are valid for Tisa programs. Another key benefit of Tisa is that its greybox specifications [11] allow service providers to encapsulate changeable implementation details by hiding them using a combination of *spec* and **refining** expressions. Thus, Tisa provides significant modularity benefits while balancing the verification needs.

References

1. Papazoglou, M.P., Georgakopoulos, D.: Service-oriented computing: Introduction. *Commun. ACM* 46(10), 24–28 (2003)
2. Christensen, E., Curbera, F., Meredith, G., Weerawarana, S.: Web services description language (WSDL) 1.1. Technical report, World Wide Web Consortium (March 2001)
3. Barth, A., Mitchell, J., Datta, A., Sundaram, S.: Privacy and utility in business processes. In: *CSF 2007*, pp. 279–294 (2007)
4. Leavens, G.T., Baker, A.L., Ruby, C.: Preliminary design of JML: a behavioral interface specification language for Java. *SIGSOFT Softw. Eng. Notes* 31(3), 1–38 (2006)
5. Barbon, F., Traverso, P., Pistore, M., Trainotti, M.: Run-time monitoring of instances and classes of web service compositions. In: *ICWS 2006*, pp. 63–71 (2006)
6. Baresi, L., Ghezzi, C., Guinea, S.: Smart monitors for composed services. In: *ICSOC 2004*, pp. 193–202 (2004)
7. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for web services. In: *POPL 2008*, pp. 261–272 (2008)
8. Kuo, D., Fekete, A., Greenfield, P., Nepal, S., Zic, J., Parastatidis, S., Webber, J.: Expressing and reasoning about service contracts in service-oriented computing. In: *ICWS 2006*, pp. 915–918 (2006)
9. Wada, H., Suzuki, J., Oba, K.: Modeling non-functional aspects in service oriented architecture. In: *IEEE International Conference on Services Computing (SCC 2006)*, pp. 222–229 (2006)
10. Parnas, D.L.: On the criteria to be used in decomposing systems into modules 15(12), 1053–1058 (1972)
11. Büchi, M., Weck, W.: The greybox approach: When blackbox specifications hide too much. Technical Report 297, Turku Center for Computer Science (August 1999)
12. Back, R.J.R., von Wright, J.: Refinement calculus, part i: sequential nondeterministic programs. In: *REX workshop*, pp. 42–66 (1990)
13. Morris, J.M.: A theoretical basis for stepwise refinement and the programming calculus. *Sci. Comput. Program.* 9(3), 287–306 (1987)
14. Edmund, M., Clarke, J., Grumberg, O., Peled, D.A.: *Model checking*. MIT Press, Cambridge (1999)
15. Shaner, S.M., Leavens, G.T., Naumann, D.A.: Modular verification of higher-order methods with mandatory calls specified by model programs. In: *OOPSLA 2007*, pp. 351–368 (2007)
16. Necula, G.C.: Proof-carrying code. In: *POPL 1997*, pp. 106–119 (1997)
17. Rajan, H., Hosamani, M.: Tisa: Towards trustworthy services in a service-oriented architecture. *IEEE Transactions on Services Computing (SOC)* 1(2) (2008)
18. Hosamani, M., Narayanappa, H., Rajan, H.: How to trust a web service monitor deployed in an untrusted environment? In: *NWESP 2007: Proceedings of the Third International Conference on Next Generation Web Services Practices*, pp. 79–84 (2007)
19. Liskov, B., Scheifler, R.: Guardians and actions: Linguistic support for robust, distributed programs. *TOPLAS* 5(3), 381–404 (1983)

20. Gordon, A.D., Pucella, R.: Validating a web service security abstraction by typing. *Formal Aspects of Computing* 17(3), 277–318 (2005)
21. Rajan, H., Leavens, G.T.: Ptolemy: A language with quantified typed events. In: Vitek, J. (ed.) ECOOP 2008. LNCS, vol. 5142, pp. 155–179. Springer, Heidelberg (2008)
22. Clifton, C., Leavens, G.T.: MiniMAO₁: Investigating the semantics of proceed. *Science of Computer Programming* 63(3), 321–374 (2006)
23. Igarashi, A., Pierce, B., Wadler, P.: Featherweight Java: A minimal core calculus for Java and GJ. In: OOPSLA 1999, pp. 132–146 (1999)
24. Flatt, M., Krishnamurthi, S., Felleisen, M.: A programmer’s reduction semantics for classes and mixins. In: *Formal Syntax and Semantics of Java*, pp. 241–269 (1999)
25. Clifton, C.: A design discipline and language features for modular reasoning in aspect-oriented programs. Technical Report 05-15, Iowa State University (Jul 2005)
26. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Information and Computation* 115(1), 38–94 (1994)
27. Rajan, H., Tao, J., Shaner, S.M., Leavens, G.T.: Reconciling trust and modularity in web services. Technical Report 08-07, Dept. of Computer Sc., Iowa State U. (July 2008)
28. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: *Proceedings of the First Symposium on Logic in Computer Science*, pp. 322–331 (1986)
29. Buchi, J.: On a decision method in restricted second order arithmetic. In: *Proc. Internat. Congr. Logic, Method. and Philos. Sci.*, pp. 1–12 (1960)
30. Barnett, M., Schulte, W.: Runtime verification of .net contracts. *Journal of Systems and Software* 65(3), 199–208 (2003)
31. Barnett, M., Schulte, W.: Spying on components: A runtime verification technique. In: *Workshop on Specification and Verification of Component-Based Systems* (2001)
32. Barnett, M., Schulte, W.: The ABCs of specification: AsmL, Behavior, and Components. *Informatica* 25(4), 517–526 (2001)
33. Wasserman, H., Blum, M.: Software reliability via run-time result-checking. *J. ACM* 44(6), 826–849 (1997)
34. Tyler, B., Soundarajan, N.: Black-box testing of grey-box behavior. In: Petrenko, A., Ulrich, A. (eds.) FATES 2003. LNCS, vol. 2931, pp. 1–14. Springer, Heidelberg (2004)
35. Bravetti, M., Zavattaro, G.: Towards a unifying theory for choreography conformance and contract compliance. In: Lumpe, M., Vanderperren, W. (eds.) SC 2007. LNCS, vol. 4829, pp. 34–50. Springer, Heidelberg (2007)
36. Acciai, L., Boreale, M.: XPi: A typed process calculus for XML messaging. *Science of Computer Programming* 71(2), 110–143 (2008)
37. Bartoletti, M., Degano, P., Ferrari, G.L.: Types and effects for secure service orchestration. In: CSFW, pp. 57–69 (2006)
38. Bartoletti, M., Degano, P., Ferrari, G.L., Zunino, R.: Semantics-based design for secure web services. *IEEE Trans. Software Eng.* 34(1), 33–49 (2008)
39. Wei, J., Singaravelu, L., Pu, C.: Guarding sensitive information streams through the jungle of composite web services. In: ICWS 2007, pp. 455–462 (2007)
40. Srivatsa, M., Iyengar, A., Mikalsen, T., Rouvellou, I., Yin, J.: An access control system for web service compositions. In: ICWS 2007, pp. 1–8 (2007)
41. Skalka, C., Wang, X.S.: Trust but verify: authorization for web services. In: SWS, pp. 47–55 (2004)
42. Skalka, C., Smith, S.F.: History effects and verification. In: Chin, W.-N. (ed.) APLAS 2004. LNCS, vol. 3302, pp. 107–128. Springer, Heidelberg (2004)
43. Biskup, J., Carminati, B., Ferrari, E., Muller, F., Wortmann, S.: Towards secure execution orders for composite web services. In: ICWS 2007, pp. 489–496 (2007)
44. Vorobiev, A., Han, J.: Specifying dynamic security properties of web service based systems. In: SKG 2006, p. 34 (2006)