

3-7-2009

Reconciling Trust and Modularity Goals in Web Services

Hridesh Rajan
Iowa State University

Jia Tao
Iowa State University, jtao510@gmail.com

Steve M. Shaner
Iowa State University, smshaner@mac.com

Gary T. Leavens
Iowa State University

Follow this and additional works at: http://lib.dr.iastate.edu/cs_techreports

 Part of the [Programming Languages and Compilers Commons](#), and the [Theory and Algorithms Commons](#)

Recommended Citation

Rajan, Hridesh; Tao, Jia; Shaner, Steve M.; and Leavens, Gary T., "Reconciling Trust and Modularity Goals in Web Services" (2009).
Computer Science Technical Reports. 275.
http://lib.dr.iastate.edu/cs_techreports/275

This Article is brought to you for free and open access by the Computer Science at Iowa State University Digital Repository. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Reconciling Trust and Modularity Goals in Web Services

Hriday Rajan, Jia Tao, Steve Shaner and Gary T. Leavens

TR #08-07b

Initial Submission: July 16, 2008.

Revised: Mar 7, 2009.

Keywords: web services, data integrity, greybox specification, refinement

CR Categories:

D.3.1 [*Programming Languages*] Formal definitions and theory — syntax and semantics

F.3.1 [*Logics and meaning of programs*] Specifying and reasoning about programs — assertion, mechanical verification, pre and postcondition

This report is an expanded version of the following paper.

Hriday Rajan, Jia Tao, Steve Shaner, and Gary T. Leavens. Tisa: A Language Design and Modular Verification Technique for Temporal Policies in Web Services, *18th European Symposium on Programming (ESOP '09)*, March 2009, York, UK. *Lecture Notes in Computer Science*, Volume 5502, Springer-Verlag, 2009. The ESOP '09 version is Copyright (c) 2009, Springer-Verlag. The authors retain the copyright on this expanded version, which is Copyright (c) 2009, Hriday Rajan, Jia Tao, Steve Shaner and Gary T. Leavens.

Department of Computer Science
226 Atanasoff Hall
Iowa State University
Ames, Iowa 50011-1041, USA

Reconciling Trust and Modularity Goals 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. We show an example 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. These new integration possibilities are valuable for constructing today’s interoperable, large-scale, complex software-intensive systems.

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,

^{*} 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.

consider a healthcare service that allows patients to make appointments and ask prescription and health-related questions from healthcare practitioners [3]. Figure 1 shows how messages are passed in this system.



Fig. 1. Overview of a healthcare service workflow, based on [3, Fig. 3].

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 [5–9].

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 de-

```

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. 2. An Example Greybox Specification

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 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 2. 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: 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], zero-knowledge proofs [17], or a hardware-based root of trust [18, 19] 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 [20] and the work of Gordon and Pucella [21]. (Furthermore, some of our descriptions of the language syntax are adapted from Ptolemy [22].) *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 [23], a variant of Featherweight Java [24] and Classic Java [25]. *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 3 and explained below. A *Tisa* program consists of zero or more declarations, and a client (see Figure 4). Declarations are either class declarations or web service declarations.

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

```

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
 $var \in \{\text{this}, \text{thisSite}\} \cup \mathcal{V}$,
 \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. 3. Abstract syntax, based on [26, Figure 3.1, 3.7].

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 4. This service contains two web-method declarations, 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`. In examples we use commas to separate method formals. 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 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 [26, 23].

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 imple-

```

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 service Doctor {
33   Queue topQ; Queue medQ; Queue lowQ;
34   int query(int pId, int msg) {
35     refining requires pId > 0 && msg >= 2
36     ensures result > 0 {
37       ticket = ticket + 1;
38       if (msg > 500) {
39         topQ.add(pId, msg, ticket + 2000)
40       } else if (msg > 250) {
41         medQ.add(pId, msg, ticket + 2000)
42       } else {
43         lowQ.add(pId, msg, ticket + 2000)
44       };
45       q.qId
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. 4. An Example Tisa Implementation

```

specification ::= servicespec*
servicespec ::= service w { wmspec* }
wmspec ::= t m (form*) { se }
form ::= t var, where var ≠ thisSite
spec ::= requires sp ensures se
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

```

Fig. 5. Syntax for Writing Specifications in Tisa

menting 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 5. In this figure, all nonterminals that are used but not defined are the same as in Figure 3. 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 3 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]. The two levels of the grammar for *se* prevent nesting of specification expressions within specification expressions.

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** \perp **ensures** sp .

An example greybox specification of the web service `Patient` appears in Figure 2. 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].

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

The language specifies histories that are sequences of web method calls. For a given *specification*, a policy can be an atomic proposition in $\mathcal{P}(\textit{specification})$; a negation of a policy or boolean combination of policies. For simplicity here we take the set of legal propositions $\mathcal{P}(\textit{specification})$ 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:

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

The constant *true* means that the service does not have any obligation. The operator **F** is read as “eventually” and **G** as “always”. Interesting temporal policies can be constructed via nesting of these temporal operators. Below we present two sample policies for our healthcare service example.

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

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`.

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

$$\begin{array}{c}
\text{(WEB METHOD CALL)} \\
\frac{\begin{array}{l} \Pi = \{\text{var}_i : \text{var } t_i \mid 1 \leq i \leq n\} \uplus \{\text{this} : \text{var } c_2\} \uplus \{\text{thisSite} : \text{var } \text{Site}\} \quad \nu = \text{frame } \rho \Pi \\ \rho = \{\text{var}_i \mapsto v_i \mid 1 \leq i \leq n\} \oplus (\text{this} \mapsto \text{loc}) \oplus (\text{thisSite} \mapsto w) \\ (\text{loc}, c_2, t \ m(t_1 \text{var}_1, \dots, t_n \text{var}_n)\{e\}) = \text{find}(w, m) \end{array}}{\langle \mathbb{E}[m(v_1, \dots, v_n)@w], J, S \rangle \hookrightarrow \langle \mathbb{E}[\text{under } e], \nu + J, S \rangle} \\
\\
\text{(REFINING)} \\
\frac{n \neq 0}{\langle \mathbb{E}[\text{refining requires } n \text{ ensures } e' \{e''\}], J, S \rangle \hookrightarrow \langle \mathbb{E}[\text{evalbody } e''e'], J, S \rangle} \\
\\
\text{(EVALBODY)} \\
\frac{\begin{array}{l} \rho = \text{envOf}(\nu) \quad \Pi = \text{tenvOf}(\nu) \quad w = \text{thisSite}(\nu) \quad t = \text{typeOf}(v, S, w) \\ \rho' = \Pi \uplus \{\text{result} : v\} \quad \Pi' = \Pi \uplus \{\text{result} : \text{var } t\} \quad \nu' = \text{frame } \rho' \Pi' \end{array}}{\langle \mathbb{E}[\text{evalbody } v e'], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[\text{under evalpost } v e'], \nu' + \nu + J, S \rangle} \\
\\
\begin{array}{cc}
\text{(EVALPOST)} & \text{(UNDER)} \\
\frac{n \neq 0}{\langle \mathbb{E}[\text{evalpost } v n], J, S \rangle \hookrightarrow \langle \mathbb{E}[v], J, S \rangle} & \langle \mathbb{E}[\text{under } v], \nu + J, S \rangle \\
& \hookrightarrow \langle \mathbb{E}[v], J, S \rangle
\end{array}
\end{array}$$

Fig. 6. Operational semantics of Tisa. Standard OO rules are presented in Section A.1.

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 [26]). 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.

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 [27] that acts on the position in the overall expression identified by the evaluation context as shown in Figure 6. Standard OO rules are presented in our technical report [28]. 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 Examples in Tisa

We have tried several examples in Tisa and they worked out wonderfully. In this section, we will discuss two other example web-services in Tisa. Each of these examples is inspired from a real-world web-service, however, we consider simplified versions of these for ease of presentation.

3.1 Streamlined Sales Tax Service

Many states in the United States enact some form of sales tax to raise revenue. E-commerce complicates the collection of such taxes due to the boundary-crossing nature of transactions on the internet. Identifying which states need to be paid and how much to pay them is a concern shared by all e-businesses. Likewise, the states want to be seen as amenable to e-business. If the states provide a suitable automated e-filing system for their sales tax, web services could be developed to distribute per-state taxes given a descriptive set of receipts by some client. A multi-state project enabling these web services is now under development [29].

For such web services demonstrating compliance to non-functional policies would be crucial. For example, the service provider may like to demonstrate that “the tax returns filed by the clients are indeed sent to the relevant state’s e-file system” to inspire client’s trust in the web service implementation. Figure 7 illustrates the greybox specifications and Figure 8 shows the implementation of the Sales Tax Service. The implementation uses the service `ZipToState` (not shown) to determine the state corresponding to the argument zip code. The tax returns are then filed to the desired state

```

1 service IAFile {
2   int fileReturn(int fedId, int amount){
3     requires fedId>0 && amount>0
4     ensures result == amount
5   }
6   int getReturn(int fedId){
7     requires fedId>0
8     ensures result >= 0
9   }
10 }
11 service FLEFile {
12   int fileReturn(int fedId, int amount){
13     requires fedId>0 && amount>0
14     ensures result == amount
15   }
16   int getReturn(int fedId){
17     requires fedId>0
18     ensures result >= 0
19   }
20 }

21 service SalesTax {
22   int process(int fedId, int zip, int amount){
23     preserve fedId > 0 && zip > 0 && amount > 0;
24     int state = getState(zip)@ZipToState;
25     if (state == 19){
26       fileReturn(fedId, amount)@IAFile
27     } else if (state == 12){
28       fileReturn(fedId, amount)@FLEFile
29     }
30     establish result == amount
31   }
32 }

```

Fig. 7. Greybox Specification of Sales Tax Service

```

1 service IAFile {
2   Hashtable db;
3   int fileReturn(int fedId, int amount) {
4     refining requires fedId>0 && amount>0
5     ensures result == amount {
6       db.set(fedId, db.get(fedId) + amount);
7       amount
8     }
9   }
10  int getReturn(int fedId){
11    refining requires fedId>0
12    ensures result >= 0 {
13      db.get(fedId)
14    }
15 }
16 service FLEFile {
17   Hashtable db;
18   RequestCache cache;
19   int fileReturn(int fedId, int amount) {
20     refining requires fedId>0 && amount>0
21     ensures result == amount {
22       cache.add(fedId, amount);
23       if(cache.size()>=10) {
24         cache.commit(db)
25       }
26     }
27   }
28   int getReturn(int fedId) {
29     refining requires fedId>0
30     ensures result >= 0 {
31       cache.commit(db);
32       db.get(fedId)
33     }
34 }
35 }

36 service SalesTax {
37   Hashtable clientFiles;
38   int process(int fedId, int zip, int amount){
39     refining preserve fedId > 0 && zip > 0 && amount > 0
40     {
41       ClientFile f = clientFiles.get(fedId);
42       f.taxAmount = f.taxAmount + amount
43     }
44     int state = getState(zip)@ZipToState;
45     int taxAmount = if (state ==19) {
46       /* IA */
47       fileReturn(fedId, amount)@IAFile
48     } else if (state == 12) {
49       /* FL */
50       fileReturn(fedId, amount)@FLEFile
51     };
52     refining establish result == amount {
53       ClientFile f = clientFiles.get(fedId);
54       f.lastState = state;
55       taxAmount
56     }
57 }
58 }

59 class ClientFile {
60   int fedId;
61   int taxAmount;
62   int lastState;
63 }

```

Fig. 8. Implementation of the Sales Tax Service

on behalf of clients by calling the web-methods `fileReturn` of that state's efile service. Using a blackbox specification, a client may not be able to tell whether the tax is paid to the correct state, whereas with the greybox specification this can be expressed. A sample policy for our sales tax example follows.

$$G(\text{process@SalesTax}) \wedge (\mathbf{XF}(\text{fileReturn@IAFile}) \vee \mathbf{XF}(\text{fileReturn@FLEFile}))$$

The policy states that whenever there is a web-method call `process` at the site `SalesTax`, there is eventually a web-method call `fileReturn` at the site `IAEFile` or `FLEFile`. Informally, for this small example this policy can be verified by just inspecting the specification of the web-method `process` in Figure 7. As we will show in Section 4.1 that policies such as the example policy above can be verified using just this greybox specification as an input. The clients of the sales tax service may use the specification to verify whether the service satisfies their desired policies. The refinement component of our verification technique, described in Section 4.2, can then be used by clients as a blackbox to check whether the service implementation refines its public specification.

This example also demonstrates the modularity benefits of Tisa’s greybox specifications. In the specification of the web-method `process` for the service `SalesTax`, the specification expressions hide much of the implementation details. For example details about how client’s accounting is kept by the web-service is not exposed in the specification. Other similar changeable implementation details such as the policy to deploy a logging mechanism or to profile the execution of the web-method to measure throughput of service requests, etc, can be easily added and removed from the service implementation. Figure 8 shows that the service implementation uses two **refining** expressions in the implementation of the web-method `process` to hide the details related to client accounting that is kept around to bill clients for tax-related services.

To illustrate the benefits of modularity, consider the implementations of the services `IAEFile` and `FLEFile`. The former uses a straightforward, write-through technique where tax returns filed are immediately committed to the database, whereas the latter uses a simple caching strategy to minimize writes to the database. Both these implementations refine similar greybox specifications, which does not expose details about how tax returns are stored internally. As a result, it becomes possible to replace the implementation of `IAEFile` with that similar to `FLEFile` (perhaps to improve efficiency) without breaking reasoning of any clients. Such replacement would not have been easy, if the entire implementation of these services were exposed for reasoning.

3.2 Web-based Photo Album

Consider an online photo album sharing application. Every user has an album containing photos. Users can see photos in their own albums and those from their friends as well. For simplicity, we design the application in a way that each photo has its own id and its owner’s id as properties of the photo. We also assume that all users are already authenticated. The web service `Album` contains all the photos and the service `Friends` has the relationship of whether two users are friends or not. This relationship is reflexive; i.e., every user is their own friend.

In practice, the photos will be saved in a persistent storage such as a file system and such file system may be organized into directories. However, to avoid complications in modeling, we represent this file system as instances of a class. We do not model directories. In the implementation, we use an instance of the `List` class to store the set of friends who can view a picture. The implementation of `List` would be standard and thus not shown in this example. Figure 9 shows the grey box specification and Figure 10 shows an implementation of this service. An example policy for such service might be:

```

1 service Album {
2   int view (int uid, int pid) {
3     preserve uid > 0 && pid > 0;
4     int pUid = photoUid(pid)@Album;
5     if (pUid == 0) { establish result == 0 }
6     else {
7       int isFr = find(pUid, uid)@Friends;
8       establish result == isFr
9     }
10  }
11  int putPhoto (int uid, int pid) {
12    requires uid > 0 && pid > 0
13    ensures ensures result >= 0
14  }
15  int photoUid(int pid) {
16    requires pid >= 0 ensures result >= 0
17  }
18 }
19 service Friends {
20  int find(int uid, int fid) {
21    requires uid > 0 && pid > 0
22    ensures result >= 0
23  }
24  int add(int uid, int fid) {
25    requires uid > 0 && pid > 0
26    ensures result == 1
27  }
28 }

```

Fig. 9. Greybox Specification of the Photo Sharing Service

$G(\text{view@Album} \wedge (\mathbf{XF} \text{find@Friends}))$

which says that when a photo is viewed, the friendship relation is always checked. However, this policy can be seen to not be followed, as it does not take into account the early return in the `view` web method for the case where the photo id (`pid`) is not found. To express a policy that takes arguments and results into account we would need a more complex policy specification language.

```

1 class Photo { int pid; int uid }
2 class List { /* ... */}
3 class FriendRel { int uid; int fid }
4 class PhotoList extends List { /* ... */
5   Photo findPhoto(int pid) {
6     if (this.empty()) {
7       Photo p = new Photo(); p.pid = 0; p }
8     else {
9       Photo p = this.car();
10      if (p.pid == pid) { p }
11      else { this.cdr().findPhoto(pid) }
12    }
13  }
14 }
15 service Album {
16   PhotoList photoList;
17   int view (int uid, int pid) {
18     refining preserve uid>0 && pid>0 {};
19     int pUid = photoUid(pid)@Album;
20     if (pUid == 0) {
21       refining establish result == 0 {0} }
22     else {
23       int isFr = find(pUid, uid)@Friends;
24       refining establish result == isFr {
25         isFr
26       } }
27   }
28   int putPhoto (int uid, int pid) {
29     refining requires uid>0 && pid>0
30     ensures result == 1 {
31       photoList.add(uid, pid); 1
32     }
33   }
34   int photoUid(int pid) {
35     refining requires pid >= 0
36     ensures result >= 0 {
37       photoList.findPhoto(pid)
38     }
39   }
40 }

41 class FriendList extends List { /* ... */
42   int findFriend(int uid, int fid) {
43     if(this.empty()){ 0 }
44     else {
45       FriendRel fr = this.car();
46       if (fr.uid == uid) {
47         if (fr.fid == fid) 1
48       }
49       else {
50         this.cdr().findFriend(uid, fid)
51       }
52     }
53   }
54   service Friends {
55     FriendList friendList;
56     int find(int uid, int fid) {
57       refining requires uid > 0 && pid > 0
58       ensures result >= 0 {
59         friendList.findFriend(uid, fid)
60       }
61     }
62     int add(int uid, int fid) {
63       refining requires uid>0 && pid>0
64       ensures result == 1 {
65         friendList.add(uid, fid); 1
66       }
67     }
68   }
69   client User{
70     add(200,250)@Friends;
71     put (200,12345)@Album;
72     view(250,12345)@Album;
73 }

```

Fig. 10. Implementation of Photo Sharing Service

4 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.

4.1 Verifying Policies

We adopt the standard automata-theoretic approach for verifying linear temporal logic formulas proposed by Vardi and Wolper [30] to verify policies in Tisa. Following Vardi and Wolper [30], 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 11 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 [30, 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 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. Finally, an edge is added from the new initial state z to the two original initial states z' and z'' .

The (SPEC EXP FSM) rule models the cases for satisfaction of precondition and postcondition. The states corresponding to precondition being true and the postcondition

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

(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)

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

(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)

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

(DEF EXP FSM)

$$\frac{NT \vdash se' \rightsquigarrow (Z', z', R', \Delta'), NT' \quad NT' \vdash se'' \rightsquigarrow (Z'', z'', R'', \Delta''), NT'' \quad Z = Z' \cup Z'' \cup z \quad R = R' \cup R'' \cup \{(z, z')\} \cup \{(z_i, z'') \mid z_i \in \mathit{final}(Z', R')\}}{NT \vdash \mathit{tvar} = se' ; se'' \rightsquigarrow (Z, z, R, \Delta), NT''}$$

(SEQ EXP FSM)

$$\frac{NT \vdash se' \rightsquigarrow (Z', z', R', \Delta'), NT' \quad NT' \vdash se'' \rightsquigarrow (Z'', z'', R'', \Delta''), NT'' \quad Z = Z' \cup Z'' \cup z \quad R = R' \cup R'' \cup \{(z, z')\} \cup \{(z_i, z'') \mid z_i \in \mathit{final}(Z', R')\}}{NT \vdash se' ; se'' \rightsquigarrow (Z, z, R, \Delta), NT''}$$

Fig. 11. Finite-state machine construction, built from expressions in a specification.

being true are z_1 and z_3 . The states z_2 and z_4 correspond to precondition being false and postcondition being false respectively. The transitions inserted in R ensure that the postcondition-related states z_3 and z_4 are only reachable from the precondition true state z_1 . The initial state of the finite-state machine generated from the expression following spec expression is z' . This state should only be reachable if both precondition and postcondition are true. Thus, an edge is added from the state z_3 to z' in R . Finally, edges are added such that the states z_1 and z_2 are reachable from the new initial state z .

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. Thus the (WEB METHOD CALL FSM 2) rule checks that the current web-method is already expanded into an FSM, and if so it uses the previously generated initial state for the web-method. To illustrate the (WEB METHOD CALL FSM 1) rule, consider a service defined by **service** w { **int** $m()$ { ... $m()$; ... }. At the call site for $m()$, suppose there is no z that NT maps the pair (w, m) to; in this case the (WEB METHOD CALL FSM 1) is used, putting the pair in the table NT' used to check the body. When the body of $m()$ is expanded, the process eventually encounters the call site again; however, this time the pair (w, m) is in the domain of NT , and so the (WEB METHOD CALL FSM 2) rule must be used,

which terminates the recursion in the process, by producing $(\{z\}, z, \{\}, \{\})$. Alternative evaluation also leads to same effective results. 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.

$$\begin{array}{c}
 \text{(SERVICE SPEC)} \\
 \frac{
 \begin{array}{c}
 NT \vdash m_1(\text{form}_{11}, \dots, \text{form}_{1k})@w \rightsquigarrow (Z_1, z_1, R_1, \Delta_1), NT_1 \\
 \dots \\
 NT_{n-1} \vdash m_n(\text{form}_{n1}, \dots, \text{form}_{nq})@w \rightsquigarrow (Z_n, z_n, R_n, \Delta_n), NT_n \\
 Z = Z_1 \cup \dots \cup Z_n \cup \{z\} \quad \Delta = \Delta_1 \uplus \dots \uplus \Delta_n \quad R = R_1 \cup \dots \cup R_n \cup \{(z, z_1), \dots, (z, z_n)\}
 \end{array}
 }{
 NT \vdash \mathbf{service} w \{t m_1(\text{form}_{11}, \dots, \text{form}_{1k})\{se_1\}, \dots, t m_n(\text{form}_{n1}, \dots, \text{form}_{nq})\{se_n\}\} \\
 \rightsquigarrow (Z, z, R, \Delta), NT_n
 }
 \end{array}$$

To verify a policy, we first use the algorithm defined in Figure 11 to compute the finite-state machine $\mathcal{F}(\mathcal{S})$. We then construct a Büchi automaton [31], $\mathcal{B}(\neg\phi(\mathcal{S}))$ for the policy $\phi(\mathcal{S})$ as shown by Vardi and Wolper [30]. Now as shown by Vardi and Wolper we compute whether $\mathcal{F}(\mathcal{S}) \cap \mathcal{B}(\neg\phi(\mathcal{S}))$ is empty. If this set is empty, we conclude that the specification \mathcal{S} satisfies the policy $\phi(\mathcal{S})$.

4.2 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 12. The operational semantics rules (REFINING), (EVALBODY) and (EVALPOST) ensure that the body of every **refining** expression obeys the specification it is refining.

4.3 Soundness of Verification Technique

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

Definition 1 (A Path for \mathcal{S}). Let \mathcal{S} be a specification and $\mathcal{F}(\mathcal{S}) = (\mathcal{Z}, z_0, R, \Delta)$ be the FSM for \mathcal{S} constructed using algorithm shown in Figure 11. A path t for \mathcal{S} is a (possibly infinite) sequence of pairs $(z_i, \Delta(z_i))$ starting with pair $(z_0, \Delta(z_0))$, where for each $i \geq 0$, $z_i \in \mathcal{Z}$ 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 $(z'_i, z'_{i+1}) \in R'$.

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

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

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

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

Fig. 12. Inference rules for proving Tisa refinement.

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 $t' = (z', \delta'_1) + \dots + (z'_n, \delta'_n) + t'_1$ 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 12. Details are contained in Section A.

Lemma 2. Given a specification S and a policy $\phi \in \Phi(S)$, the automaton $\mathcal{F}(S) \cap \mathcal{B}(-\phi)$ 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 [30, pp. 4,6]. Details are contained in Section A.

Theorem 1. Let S be a specification, ϕ be a policy in $\Phi(S)$, and P be a program. Let ϕ be satisfied by the specification S and P be a refinement of S (as defined in Figure 12). Then the policy ϕ is satisfied by the program P .

Proof Sketch: The proof follows from lemmas 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 S . Thus, ϕ , which is written over $\mathcal{P}(S)$, is also satisfied for P .

5 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 [32, 33] have considered using greybox specifications written in AsmL [34] for verifying contracts for .NET framework. Wasserman and Blum [35] also use a restricted form of greybox specifications for verification. Tyler and Soundarajan [36] 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 [37] 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 [38], Kuo *et al.* [8], Baresi *et al.* [6], Barbon *et al.* [5], Mahbub and Spanoudakis [39], etc.

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.

Acciai and Boreale attempt a similar typed technique with their work on XPi, a process calculus for XML messaging systems. Their system guarantees runtime safety for message size and structure in well-typed services. In contrast, our language provides sound refinement and modular specification for web services.

The focus of Kuo *et al.*’s approach is on facilitating a more concise representation of the message exchange protocols as Boolean formula associated with each exchanged message, which in turn helps verify whether a given message exchange is legal. On the other end of the spectrum are approaches to validate the functional and non-functional

requirements of a web service such as by Baresi *et al.* [6], Barbon *et al.* [5], Mahbub and Spanoudakis [39], etc, which use dynamic monitoring to ensure that a service-oriented architecture is satisfying its requirements. These techniques rely on monitoring the functional interface, often during service composition, to determine conformance of a web service to its requirement. Non-functional requirements such as observable web service calls are not addressed.

Wada *et al.* proposed a UML profile to graphically model non-functional aspects in SOA so that they are incorporated in the development phase [9]. This UML profile includes certain key model elements of service oriented architecture such as *service, message exchange, message, connector and filter*. This model driven development (MDD) paradigm for addressing non-functional concerns such as security and integrity in the service oriented architecture is an encouraging step for developing a secure service oriented architecture, however, it does not help with verification of observable web service calls for existing service-oriented architectures.

Another approach towards achieving trust is Aglet [40]. An Aglet is a Java object with a code component and a data component. The key idea here is to use these mobile agents to preserve privacy. An Aglet consists of two distinct parts: the Aglet core and the Aglet proxy. The Aglet core contains all the internal variables and methods. It provides interfaces through which the environment can make use of the Aglet or vice versa. The core is encapsulated with an Aglet proxy which acts as a shield against any attempt to directly access the private variables and methods of the aglet. This Aglet proxy can be programmed to enforce local privacy requirements on the site of the remote entity. Aglets are deployed into Aglet servers, which enforce the requirement of the security model. A key problem with Aglets is that the integrity of Aglets depends on the integrity of Aglet servers. This approach does not provide a technique to guarantee such integrity in an untrustworthy environment. Furthermore, this approach requires service implementations (source) to change to use the Aglets instead of the original objects.

A very basic architecture for addressing Security aspects in Web Service composition has been proposed by Charfi *et al.* [41]. The authors considered security attributes as boolean propositions and classified three classes of Security constraints. They are the general or global security constraints (have to be satisfied by every service component). Component specific security constraints and Compatibility security constraints ensure that the mutual obligations of each of the participating services in a composition is satisfied. However, their approach is restricted only to satisfy basic security assertions given the information during static composition, and does not address how composite service security attributes can be verified dynamically at an instance level. Moreover, there is a need for a verifiable composition procedure that can guarantee both static and dynamic privacy and security aspects for the generated composition.

Bartoletti *et al.* [42] 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 [43] as possible future work for Tisa.

Another approach [44] 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.* [45] 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 [46] 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 [47] 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 [48, 49] either do not have a formal model supporting them or are tightly coupled with implementations.

Language-based Information flow techniques. There is significant body of work on language-based approaches to analyzing information flow (cf. [50–68] and [69] for a survey). These techniques statically analyze and/or type-check code for secure information flow and are quite useful at the time of development. The major disadvantage of all these approaches is that they require source code. Thus, they cannot be applied transparently to already deployed applications that are only available as binaries. These techniques are also not applicable to scenarios where service provider’s implementation is not accessible (primarily due to intellectual property issues).

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 [5–9], 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) (2003) 24–28
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'07. 279–294
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) (2006) 1–38
5. Barbon, F., Traverso, P., Pistore, M., Trainotti, M.: Run-time monitoring of instances and classes of web service compositions. In: ICWS '06. 63–71
6. Baresi, L., Ghezzi, C., Guinea, S.: Smart monitors for composed services. In: ICSOC '04. 193–202
7. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for web services. In: POPL '08. 261–272
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 '06. 915–918
9. Wada, H., Suzuki, J., Oba, K.: Modeling non-functional aspects in service oriented architecture. In: IEEE International Conference on Services Computing (SCC'06). (2006) 222–229
10. Parnas, D.L.: On the criteria to be used in decomposing systems into modules. *Communications of the ACM* **15**(12) (December 1972) 1053–8
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. (1990) 42–66
13. Morris, J.M.: A theoretical basis for stepwise refinement and the programming calculus. *Sci. Comput. Program.* **9**(3) (1987) 287–306
14. Edmund M. Clarke, J., Grumberg, O., Peled, D.A.: Model checking. MIT Press, Cambridge, MA, USA (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 '07. 351–368
16. Necula, G.C.: Proof-carrying code. In: POPL '97. 106–119
17. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in np have zero-knowledge proof systems. *J. ACM* **38**(3) (1991) 690–728
18. Rajan, H., Hosamani, M.: Tisa: Towards trustworthy services in a service-oriented architecture. *IEEE Transactions on Services Computing (SOC)* **1**(2) (2008)
19. Hosamani, M., Narayanappa, H., Rajan, H.: How to trust a web service monitor deployed in an untrusted environment? In: NWESP '07: Proceedings of the Third International Conference on Next Generation Web Services Practices. (2007) 79–84
20. Liskov, B., Scheifler, R.: Guardians and actions: Linguistic support for robust, distributed programs. *TOPLAS* **5**(3) (July 1983) 381–404
21. Gordon, A.D., Pucella, R.: Validating a web service security abstraction by typing. *Formal Aspects of Computing* **17**(3) (2005) 277–318
22. Rajan, H., Leavens, G.T.: Ptolemy: A language with quantified typed events. In: 22nd European Conference on Object-oriented Programming (ECOOP 2008). (July 2008)
23. Clifton, C., Leavens, G.T.: MiniMAO₁: Investigating the semantics of proceed. *Science of Computer Programming* **63**(3) (2006) 321–374
24. Igarashi, A., Pierce, B., Wadler, P.: Featherweight Java: A minimal core calculus for Java and GJ. In: OOPSLA '99. 132–146

25. Flatt, M., Krishnamurthi, S., Felleisen, M.: A programmer's reduction semantics for classes and mixins. In: *Formal Syntax and Semantics of Java*. Springer-Verlag (1999) 241–269
26. Clifton, C.: A design discipline and language features for modular reasoning in aspect-oriented programs. Technical Report 05-15, Iowa State University (Jul 2005)
27. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Information and Computation* **115**(1) (Nov 1994) 38–94
28. 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)
29. : Streamlined sales tax project. <http://www.streamlinedsalestax.org/>
30. Vardi, M.Y., Wolper, P.: An automata-theoretic approach to automatic program verification. In: *Proceedings of the First Symposium on Logic in Computer Science*. (1986) 322–331
31. Buchi, J.: On a decision method in restricted second order arithmetic. *Proc. Internat. Congr. Logic, Method. and Philos. Sci* (1960) 1–12
32. Barnett, M., Schulte, W.: Runtime verification of .net contracts. *Journal of Systems and Software* **65**(3) (March 2003) 199–208
33. Barnett, M., Schulte, W.: Spying on components: A runtime verification technique. In: *Workshop on Specification and Verification of Component-Based Systems*. (2001)
34. Barnett, M., Schulte, W.: The ABCs of specification: AsmL, Behavior, and Components. *Informatica* **25**(4) (November 2001) 517–526
35. Wasserman, H., Blum, M.: Software reliability via run-time result-checking. *J. ACM* **44**(6) (1997) 826–849
36. Tyler, B., Soundarajan, N.: Black-box testing of grey-box behavior. In: *FATES '03*, 1–14.
37. Bravetti, M., Zavattaro, G.: Towards a unifying theory for choreography conformance and contract compliance. In: *Software Composition*. (2007) 34–50
38. Acciai, L., Boreale, M.: XPi: A typed process calculus for XML messaging. *Science of Computer Programming* **71**(2) (2008) 110–143
39. Mahbub, K., Spanoudakis, G.: Run-time monitoring of requirements for systems composed of web-services: Initial implementation and evaluation experience. In: *ICWS '05*. 257–265
40. Rezgui, A., Ouzzani, M., Bouguettaya, A., Medjahed, B.: Preserving privacy in web services. In: *WIDM '02*. 56–62
41. Charfi, A., Mezini, M.: Using aspects for security engineering of web service compositions. In: *ICWS '05*. 59–66
42. Bartoletti, M., Degano, P., Ferrari, G.L.: Types and effects for secure service orchestration. In: *CSFW*. (2006) 57–69
43. Bartoletti, M., Degano, P., Ferrari, G.L., Zunino, R.: Semantics-based design for secure web services. *IEEE Trans. Software Eng.* **34**(1) (2008) 33–49
44. Wei, J., Singaravelu, L., Pu, C.: Guarding sensitive information streams through the jungle of composite web services. In: *ICWS '07*. 455–462
45. Srivatsa, M., Iyengar, A., Mikalsen, T., Rouvellou, I., Yin, J.: An access control system for web service compositions. In: *ICWS '07*. 1–8
46. Skalka, C., Wang, X.S.: Trust but verify: authorization for web services. In: *SWS*. (2004) 47–55
47. Skalka, C., Smith, S.F.: History effects and verification. In: *APLAS*. (2004) 107–128
48. Biskup, J., Carminati, B., Ferrari, E., Muller, F., Wortmann, S.: Towards secure execution orders for composite web services. In: *ICWS '07*. 489–496
49. Vorobiev, A., Han, J.: Specifying dynamic security properties of web service based systems. In: *SKG '06*. 34–34
50. Bell, D., LaPadula, L.: *Secure Computer Systems: Mathematical Foundations*. DTIC Research Report AD0770768 (1973)

51. Chong, S., Liu, J., Myers, A.C., Qi, X., Vikram, K., Zheng, L., Zheng, X.: Secure web applications via automatic partitioning. In: 21st ACM Symposium on Operating Systems Principles (SOSP'07). (October) 31–44
52. Chong, S., Vikram, K., Myers, A.C.: Sif: Enforcing confidentiality and integrity in web applications. In: USENIX Security Symposium 2007. (August 2007) 1–16
53. Chess, B., McGraw, G.: Static analysis for security. *Security & Privacy Magazine, IEEE* **2**(6) (2004) 76–79
54. Livshits, V.B., Lam, M.S.: Finding security vulnerabilities in java applications with static analysis. In: SSYM'05: Proceedings of the 14th conference on USENIX Security Symposium, Berkeley, CA, USA, USENIX Association (2005) 18–18
55. Evans, D., Larochelle, D.: Improving security using extensible lightweight static analysis. *IEEE Softw.* **19**(1) (2002) 42–51
56. Evans, D., Twyman, A.: Flexible policy-directed code safety. In: IEEE Symposium on Security and Privacy. (1999) 32–45
57. Myers, A.C.: Jflow: practical mostly-static information flow control. In: POPL '99: Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, New York, NY, USA, ACM (1999) 228–241
58. Fournet, C., Rezk, T.: Cryptographically sound implementations for typed information-flow security. In: POPL '08: Proceedings of the 35th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages, New York, NY, USA, ACM (2008) 323–335
59. Smith, G., Volpano, D.: Secure information flow in a multi-threaded imperative language. In: POPL '98: Proceedings of the 25th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, New York, NY, USA, ACM (1998) 355–364
60. Volpano, D., Smith, G.: Verifying secrets and relative secrecy. In: POPL '00: Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, New York, NY, USA, ACM (2000) 268–276
61. Hristova, K., Rothamel, T., Liu, Y.A., Stoller, S.D.: Efficient type inference for secure information flow. In: PLAS '06: Proceedings of the 2006 workshop on Programming languages and analysis for security, New York, NY, USA, ACM (2006) 85–94
62. Echahed, R., Prost, F., Serwe, W.: Statically assuring secrecy for dynamic concurrent processes. In: PPDP '03: Proceedings of the 5th ACM SIGPLAN international conference on Principles and practice of declarative programming, New York, NY, USA, ACM (2003) 91–101
63. Riely, J., Hennessy, M.: Trust and partial typing in open systems of mobile agents. In: POPL '99: Proceedings of the 26th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, New York, NY, USA, ACM (1999) 93–104
64. Barthe, G., Nieto, L.P.: Formally verifying information flow type systems for concurrent and thread systems. In: FMSE '04: Proceedings of the 2004 ACM workshop on Formal methods in security engineering, New York, NY, USA, ACM (2004) 13–22
65. Hennessy, M., Riely, J.: Information flow vs. resource access in the asynchronous pi-calculus. *ACM Trans. Program. Lang. Syst.* **24**(5) (2002) 566–591
66. Vitek, J., Bokowski, B.: Confined types. In: OOPSLA '99: Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications, New York, NY, USA, ACM (1999) 82–96
67. Abadi, M.: Secrecy by typing in security protocols. *J. ACM* **46**(5) (1999) 749–786
68. Myers, A.C., Liskov, B.: Protecting privacy using the decentralized label model. *ACM Trans. Softw. Eng. Methodol.* **9**(4) (2000) 410–442
69. Sabelfeld, A., Myers, A.: Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* **21**(1) (Jan 2003) 5–19
70. Schmidt, D.A.: *The Structure of Typed Programming Languages*. Foundations of Computing Series. MIT Press, MA (1994)

71. Choueka, Y.: Theories of automata on ω -tapes: A simplified approach. *Journal of Computer and System Sciences* **8**(2) (1974) 117–141
72. Emerson, E.A., Lei, C.L.: Modalities for model checking: branching time logic strikes back. *Sci. Comput. Program.* **8**(3) (1987) 275–306

A Appendix

A.1 Omitted Details about Operational Semantics

This section presents the omitted details about the small step operational semantics for Tisa programs. Small steps are taken in the semantics to transition from one configuration to another. These configurations appear in Figure 13. A configuration contains an expression (e), a 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.

$$\text{thisSite}(\mathbf{frame} \rho \Pi) = \rho(\mathbf{thisSite})$$

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 type environment Π (see Figure 18) is not used by the operational semantics, but only in the type soundness proof.

As is usual [27] the semantics is presented as a set of evaluation contexts \mathbb{E} and an one-step reduction relation that acts on the position in the overall expression identified by the evaluation context. Figure 13 also defines the evaluation contexts and the order of evaluation for Tisa. The language uses a strict leftmost, innermost evaluation policy, which thus uses call-by-value, except for the short-circuit operators $\&\&$ and $||$, and the if-expression.

The operational semantics make implicit use of a fixed (global) class table, CT , that maps class names (including the subset of web service names) to declarations (of classes or services). It also uses a fixed instance table, IT , that maps web service names to locations; these are the locations of fixed instances that act as receiver objects for web-method calls. Both of these tables are implicitly used by various auxiliary functions.

Thus to define an initial configuration, we first have to describe how CT and IT are defined for a given program. To explain this initialization, consider the program given in Figure 14, where without loss of generality, all class declarations are placed in front of all service declarations, and $n \geq 0$, $k \geq 0$. For the program in Figure 14(left), we define the class table to map each class to its class declaration and each service declaration to the declaration of a class with the same name that inherits from `Site` as shown in Figure 14(right). For this program, we define the instance table by:

$$\begin{aligned} IT(w_{n+j}) &= loc_{n+j}, \text{ for } j \in \{1, \dots, k\} \\ IT(w_{cl}) &= \mathbf{null}. \end{aligned}$$

Added Syntax:

$$e ::= loc \mid \mathbf{under} \ e \mid \mathbf{evalbody} \ e \ e \mid \mathbf{evalpost} \ e \ e$$

| NullPointerException | ClassCastException | SpecException
where $loc \in \mathcal{L}$, a set of locations

Domains:

$\Gamma ::= \langle e, J, S \rangle$	“Configurations”
$J ::= \nu + J \mid \bullet$	“Stacks”
$\nu ::= \mathbf{frame} \ \rho \ \Pi$	“Frames”
$\rho ::= \{var_k : v_k\}_{k \in K}$,	“Environments”
where K is finite, $K \subseteq I$	
$v ::= n \mid w \mid loc \mid \mathbf{null}$	“(Expressible) Values”
$S ::= \{(w_k, loc_k) \mapsto sv_k\}_{k \in K}$,	“Stores”
where K is finite	
$sv ::= o$	“Storable Values”
$o ::= [c, F]$	“Object Records”
$F ::= \{f_k \mapsto v_k\}_{k \in K}$,	“Field Maps”
where K is finite	

Evaluation contexts:

$$\mathbb{E} ::= - \mid \mathbb{E} == e \mid v == \mathbb{E} \mid \mathbb{E} != e \mid v != \mathbb{E} \mid \mathbb{E} > e \mid v > \mathbb{E} \mid \mathbb{E} < e \mid v < \mathbb{E} \mid \mathbb{E} >= e \mid v >= \mathbb{E} \mid \mathbb{E} <= e \mid v <= \mathbb{E} \mid \mathbb{E} + e \mid v + \mathbb{E} \mid \mathbb{E} - e \mid v - \mathbb{E} \mid \mathbb{E} * e \mid v * \mathbb{E} \mid !\mathbb{E} \mid \mathbb{E} \ \&\& \ e \mid \mathbb{E} \ \&\mid \ e \mid \mathbf{isNull}(\mathbb{E}) \mid \mathbf{if}(\mathbb{E}) \ e \ \mathbf{else} \ e \mid \mathbb{E}.m(e\dots) \mid v.m(v\dots \mathbb{E}e\dots) \mid \mathbb{E}.f \mid \mathbb{E}.f=e \mid v.f=\mathbb{E} \mid \mathbb{E};e \mid \mathbf{cast} \ c \ \mathbb{E} \mid t \ \mathbf{var}=\mathbb{E}; \ e \mid m(v\dots \mathbb{E}e\dots)e \mid m(v\dots)\mathbb{E} \mid \mathbf{refining} \ \mathbf{requires} \ \mathbb{E} \ \mathbf{ensures} \ e \ \{ \ e \} \mid \mathbf{evalbody} \ \mathbb{E} \ e \mid \mathbf{evalpost} \ v \ \mathbb{E} \mid \mathbf{under} \ \mathbb{E}$$

Fig. 13. Added syntax, domains, and evaluation contexts used in the semantics, based on [26].

<pre> class c₁ extends d₁ {field*₁ meth*₁} ... class c_n extends d_n {field*_n meth*_n} service w_{n+1} {field*_{n+1} meth*_{n+1}} ... service w_{n+k} {field*_{n+k} meth*_{n+k}} client w_{cl} {e_{cl}} </pre>	$CT(\text{Site}) = \mathbf{class} \ \text{Site} \ \mathbf{extends} \ \text{Object}\{\},$ $CT(c_i) = \mathbf{class} \ c_i \ \mathbf{extends} \ d_i \ \{field*_i \ meth*_i\},$ <p style="text-align: center;">where $i \in \{1, \dots, n\}$</p> $CT(w_{n+j}) = \mathbf{class} \ w_{n+j} \ \mathbf{extends} \ \text{Site}\{$ <p style="text-align: center;">$field*_{n+j} \ meth*_{n+j}\},$</p> <p style="text-align: center;">where $j \in \{1, \dots, k\}$</p>
---	--

Fig. 14. A schematic example of a program and corresponding class table. Here $n \geq 0$.

With these definitions, we can give a well-defined initial configuration for the program in Figure 14, which is $\langle e, J_{init}, S_{init} \rangle$, where the initial stack and store are:

$$J_{init} = \mathbf{frame} \ \rho_{init} \ \Pi_{init} + \bullet$$

$$\rho_{init} = \{\mathbf{thisSite} : w_{cl}\}$$

$$\Pi_{init} = \{\mathbf{thisSite} : \text{Site}\}$$

$$S_{init} = \{(w_{n+j}, loc_{n+j}) \mapsto \mathit{initialObj}(w_{n+j})\}_{j \in \{1, \dots, n\}}$$

where $\mathit{initialObj}(c) = [c. \mathit{initialFields}(c)]$
and $\mathit{initialFields}(c) = \{f \mapsto \mathbf{null} \mid f \in \mathit{dom}(\mathit{fieldsOf}(c))\}$.

Figure 6 and Figure 15 present the operational semantics of Tisa. In these rules all of the hypotheses are really side conditions and side definitions for use in the rule. Several of the rules manipulate type information; this information is not used by the semantics, but is kept for the type soundness proof.

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

$$\begin{array}{c}
\text{(GR)} \\
\frac{n = (\mathbf{if } v_1 > v_2 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[v_1 > v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle} \\
\\
\begin{array}{ccc}
\text{(LE)} & \text{(GREQ)} & \text{(LEEQ)} \\
\frac{n = (\mathbf{if } v_1 < v_2 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[v_1 < v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle} & \frac{n = (\mathbf{if } v_1 >= v_2 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[v_1 >= v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle} & \frac{n = (\mathbf{if } v_1 <= v_2 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[v_1 <= v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle}
\end{array} \\
\\
\begin{array}{ccc}
\text{(EQEQ)} & \text{(NEQ)} & \text{(PLUS)} \\
\frac{n = (\mathbf{if } v_1 = v_2 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[v_1 = v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle} & \frac{n = (\mathbf{if } v_1 = v_2 \mathbf{ then } 0 \mathbf{ else } 1)}{\langle \mathbb{E}[v_1 != v_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle} & \frac{n_3 = n_1 + n_2}{\langle \mathbb{E}[n_1 + n_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n_3], J, S \rangle}
\end{array} \\
\\
\begin{array}{cccc}
\text{(MINUS)} & \text{(TIMES)} & \text{(NOT)} & \text{(ANDTRUE)} \\
\frac{n_3 = n_1 - n_2}{\langle \mathbb{E}[n_1 - n_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n_3], J, S \rangle} & \frac{n_3 = n_1 \times n_2}{\langle \mathbb{E}[n_1 * n_2], J, S \rangle \hookrightarrow \langle \mathbb{E}[n_3], J, S \rangle} & \frac{n_2 = (\mathbf{if } n_1 = 0 \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[!n_1], J, S \rangle \hookrightarrow \langle \mathbb{E}[n_2], J, S \rangle} & \frac{n_1 \neq 0}{\langle \mathbb{E}[n_1 \&\&e], J, S \rangle \hookrightarrow \langle \mathbb{E}[e], J, S \rangle}
\end{array} \\
\\
\begin{array}{cccc}
\text{(ANDFALSE)} & \text{(ORFALSE)} & \text{(ORTRUE)} & \text{(ISNULL)} \\
\frac{}{\langle \mathbb{E}[0 \&\&e], J, S \rangle \hookrightarrow \langle \mathbb{E}[0], J, S \rangle} & \frac{}{\langle \mathbb{E}[0 | e], J, S \rangle \hookrightarrow \langle \mathbb{E}[e], J, S \rangle} & \frac{n_1 \neq 0}{\langle \mathbb{E}[n_1 | e], J, S \rangle \hookrightarrow \langle \mathbb{E}[n_1], J, S \rangle} & \frac{n = (\mathbf{if } v = \mathbf{null} \mathbf{ then } 1 \mathbf{ else } 0)}{\langle \mathbb{E}[\mathbf{isNull}(v)], J, S \rangle \hookrightarrow \langle \mathbb{E}[n], J, S \rangle}
\end{array} \\
\\
\text{(IFTRUE)} \quad \frac{n \neq 0}{\langle \mathbb{E}[\mathbf{if}(n) \{e_2\} \mathbf{ else } \{e_3\}], J, S \rangle \hookrightarrow \langle \mathbb{E}[e_2], J, S \rangle} \quad \text{(IFFALSE)} \quad \langle \mathbb{E}[\mathbf{if}(0) \{e_2\} \mathbf{ else } \{e_3\}], J, S \rangle \hookrightarrow \langle \mathbb{E}[e_3], J, S \rangle \\
\\
\text{(NEW)} \\
\frac{w = \mathbf{thisSite}(\nu) \quad (w, loc) \notin \mathit{dom}(S) \quad S' = S \oplus ((w, loc) \mapsto [c. \{f \mapsto \mathbf{null} \mid f \in \mathit{dom}(\mathit{fieldsOf}(c))\}])}{\langle \mathbb{E}[\mathbf{new } c()], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[loc], \nu + J, S' \rangle} \\
\\
\text{(VAR)} \\
\frac{\rho = \mathit{envOf}(\nu) \quad v = \rho(\mathit{var})}{\langle \mathbb{E}[\mathit{var}], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[v], \nu + J, S \rangle} \\
\\
\text{(CALL)} \\
\frac{w = \mathbf{thisSite}(\nu) \quad [c.F] = S(w, loc) \quad (c_2, t \ m(t_1 \mathit{var}_1, \dots, t_n \mathit{var}_n) \{e\}) = \mathit{mbody}(p, c, m) \quad \rho' = \{\mathit{var}_i \mapsto v_i \mid 1 \leq i \leq n\} \oplus (\mathbf{this} \mapsto loc) \oplus (\mathbf{thisSite} \mapsto w) \quad \Pi' = \{\mathit{var}_i : \mathbf{var } t_i \mid 1 \leq i \leq n\} \uplus \{\mathbf{this} : \mathbf{var } c_2\} \uplus \{\mathbf{thisSite} : \mathbf{var } \mathit{Site}\} \quad \nu' = \mathbf{frame } \rho' \ \Pi'}{\langle \mathbb{E}[loc.m(v_1, \dots, v_n)], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e], \nu' + \nu + J, S \rangle} \\
\\
\text{(DEF)} \\
\frac{\rho' = \rho \oplus (\mathit{var} \mapsto v) \quad \Pi = \mathit{tenvOf}(\nu) \quad \Pi' = \Pi \uplus \{\mathit{var} : \mathbf{var } t\} \quad \nu' = \mathbf{lexframe } \rho' \ \Pi'}{\langle \mathbb{E}[t \ \mathit{var} = v; e], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } e], \nu' + \nu + J, S \rangle} \\
\\
\text{(SKIP)} \quad \langle \mathbb{E}[v; e], J, S \rangle \hookrightarrow \langle \mathbb{E}[e], J, S \rangle \quad \text{(CAST)} \quad \frac{w = \mathbf{thisSite}(\nu) \quad [c'.F] = S(w, loc) \quad c' \preceq c}{\langle \mathbb{E}[\mathbf{cast } c \ \mathit{loc}], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[loc], \nu + J, S \rangle} \\
\\
\text{(NCAST)} \quad \frac{}{\langle \mathbb{E}[\mathbf{cast } t \ \mathbf{null}], J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{null}], J, S \rangle} \quad \text{(GET)} \quad \frac{w = \mathbf{thisSite}(\nu) \quad [c.F] = S(w, loc) \quad v = F(f)}{\langle \mathbb{E}[loc.f], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[v], \nu + J, S \rangle} \\
\\
\text{(SET)} \quad \frac{w = \mathbf{thisSite}(\nu) \quad [c.F] = S(w, loc) \quad S' = S \oplus ((w, loc) \mapsto [c.F \oplus (f \mapsto v)])}{\langle \mathbb{E}[loc.f = v], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[v], \nu + J, S' \rangle}
\end{array}$$

Fig. 15. Operational semantics for the OO expressions in Tisa, based on [26].

The (NEW) rule says that the store is updated to map a fresh location to an object of the given class that has each of its fields set to null. This rule (and others) uses \oplus as an overriding operator for finite functions. That is, if $S' = S \oplus ((w, loc) \mapsto v)$, then $S'(w', loc') = v$ if $(w', loc') = (w, loc)$ and otherwise $S'(w', loc') = S(w', loc')$. The *fieldsOf* function uses the class table to determine the list of field declarations for a given class (and its superclasses), considered as a mapping from field names to their types.

In the (VAR) rule, *envOf*(ν) returns the environment from the current frame ν , ignoring any other information in ν .

$$envOf(\mathbf{frame} \rho \Pi) = \rho$$

Thus the (VAR) rule says that the value of a variable, including **this**, is simply looked up in the environment of the current frame. The (CALL) rule implements dynamic dispatch by looking up the method m starting from the dynamic class (c) of the receiver object (loc), looking in superclasses if necessary, using the auxiliary function *mbody* (see Figure 16). The body is executed in a **frame** with an environment that binds the methods formals, including **this**, to the actual parameters. Since methods do not nest, and since expressions access object fields by starting from an explicit object there is no other context available to a method.

$$\begin{array}{c}
 \text{(MBODY)} \\
 \hline
 CT(c) = \mathbf{class} \ c \ \mathbf{extends} \ d \ \{ \mathit{field}^* \ \mathit{meth}_1, \dots, \mathit{meth}_k \} \quad \mathit{meth}_i = t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\} \\
 \hline
 \mathit{mbody}(c, m) = (c, t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\}) \\
 \\
 \text{(MBODY)} \\
 \hline
 CT(c) = \mathbf{class} \ c \ \mathbf{extends} \ d \ \{ \mathit{field}^* \ \mathit{meth}_1, \dots, \mathit{meth}_k \} \\
 \quad \#i \in \{1, \dots, k\}. \mathit{meth}_i = t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\} \\
 \hline
 \mathit{mbody}(d, m) = (d', t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\}) \\
 \hline
 \mathit{mbody}(c, m) = (d', t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\}) \\
 \\
 \text{(FIND)} \\
 \hline
 IT(w) = loc \quad \mathit{mbody}(w, m) = (c', t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\}) \\
 \hline
 \mathit{find}(w, m) = (loc, c', t \ m(\mathit{var}_1, \dots, \mathit{var}_n)\{e\})
 \end{array}$$

Fig. 16. Auxiliary functions for looking up and finding methods.

Note that **under** e is used in the resulting configuration for the (WEB METHOD CALL) rule. This expression is used whenever a new frame is pushed on the stack, to record that the stack should be popped when the evaluation of e is finished. The (UNDER) rule pops the stack when evaluation of its subexpression is finished.

The (DEF) rule allows for local definitions. It is similar to **let** in other languages, but with a more C++ and Java-like syntax. It simply binds the variable given to the value in an extended environment. (Note that *tenvOf*(ν) is the type environment from the frame ν .) Since a new frame is pushed on the stack, the body, e , is evaluated inside

an “under” expression, which pops the stack and the principal stack when e is finished. The (SKIP) rule for sequence expressions is similar, but no new frame is needed.

The (CAST) rule simply checks that the dynamic class of the object is a subtype of the type given in the expression. The (NCAST) rule allows **null** to be cast to any type.

The (GET) and (SET) rules are standard. The value of a field assignment is the value being assigned.

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** $\text{evalpost } v e'$, 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 , which is defined as follows.

$$\begin{aligned} \text{typeOf}(n, S, w) &= \mathbf{int} \\ \text{typeOf}(loc, S, w) &= c, \text{ if } S(w, loc) = [c.F] \text{ for some } F \\ \text{typeOf}(\mathbf{null}, S, w) &= \mathbf{Object} \\ \text{typeOf}(w', S, w) &= w' \end{aligned}$$

Finally, the (EVALPOST) rule checks that the postcondition is non-zero (true) and uses the body’s value as the value of the expression.

The operational semantics rules that result in exceptions are given in Figure 17. These treat some uses of null values and bad casts as exceptions, following Java. Encountering one of these exceptions does not make the semantics be “stuck” and hence the situations that lead to these exceptions are not considered to be type errors. However, all of the resulting configurations are terminal.

$$\begin{aligned} & \text{(NCALL)} \\ & \langle \mathbb{E}[\mathbf{null}.m(v_1, \dots, v_n)], J, S \rangle \hookrightarrow \langle \mathbf{NullPointerException}, \bullet, S, W \rangle \\ & \text{(NGET)} \\ & \langle \mathbb{E}[\mathbf{null}.f], J, S \rangle \hookrightarrow \langle \mathbf{NullPointerException}, \bullet, S \rangle \\ & \text{(NSET)} \\ & \langle \mathbb{E}[\mathbf{null}.f = v], J, S \rangle \hookrightarrow \langle \mathbf{NullPointerException}, \bullet, S \rangle \\ & \text{(XCAST)} \\ & \frac{w = \text{thisSite}(v) \quad [c.F] = S((w, loc)) \quad c \not\leq t}{\langle \mathbb{E}[\mathbf{cast } t \text{ loc}], \nu + J, S \rangle \hookrightarrow \langle \mathbf{ClassCastException}, \bullet, S \rangle} \\ & \text{(FPRE)} \\ & \frac{n = 0}{\langle \mathbb{E}[\mathbf{refining } \textit{requires } n \textit{ ensure } e' \{e''\}], J, S \rangle \hookrightarrow \langle \mathbf{SpecException}, \bullet, S \rangle} \\ & \text{(FPOST)} \\ & \frac{n = 0}{\langle \mathbb{E}[\mathbf{evalpost } v n], J, S \rangle \hookrightarrow \langle \mathbf{SpecException}, \bullet, S \rangle} \end{aligned}$$

Fig. 17. Operational semantics of expressions that produce exceptions, based on [26].

A.2 Type Checking

Type checking uses the type attributes defined in Figure 18. (These use some of the notation and ideas from Schmidt’s book [70].)

$\theta ::=$	“type attributes”
OK	“program/top-level decl.”
OK in c	“meth./web-meth.”
var t	“var/formal/field”
exp τ	“expression”
specFor t	“specification”
$\tau ::= t \mid \perp$	“type expressions”
$\pi, \Pi ::= \{I : \theta_I\}_{I \in K}$,	“type environments”
where K is finite, $K \subseteq (\mathcal{L} \cup \{\mathbf{this}\}) \cup \mathcal{V}$	

Fig. 18. Type attributes.

(CHECK PROGRAM)	
$(\forall i \in \{1..n\} :: \vdash decl_i : \text{OK})$	$\vdash e : \mathbf{exp\ int}$
<hr/>	
$\vdash decl_1 \dots decl_n \ \mathbf{client\ w\ \{e\}} : \text{OK}$	
(CHECK CLASS)	
$(\forall i \in \{1..n\} :: isType(t_i) \wedge f_i \notin dom(fieldsOf(d)))$	$(\forall j \in \{1..m\} :: \vdash meth_j \text{OK in } c)$
$isClass(d)$	
<hr/>	
$\vdash \mathbf{class\ } c \ \mathbf{extends\ } d \{ t_1 \ f_1; \dots \ t_n \ f_n; \ \mathit{meth}_1 \ \dots \ \mathit{meth}_m \} : \text{OK}$	
(CHECK METHOD)	
$(\mathbf{class\ } c \ \mathbf{extends\ } d \{ \dots \}) \in CT$	
$override(m, d, t_1 \times \dots \times t_n \rightarrow t)$	$(\forall i \in \{1..n\} :: isType(t_i))$
$isType(t)$	$\Pi' = \{var_1 : \mathbf{var\ } t_1, \dots, var_n : \mathbf{var\ } t_n, \mathbf{this} : \mathbf{var\ } c, \mathbf{thisSite} : \mathbf{var\ Site}\}$
$\Pi' \vdash e : \mathbf{exp\ } \tau' \quad \tau' \preceq t$	
<hr/>	
$\vdash t \ m(t_1 \ var_1, \dots, t_n \ var_n) \{e\} : \text{OK in } c$	

Fig. 19. Type-checking rules for standard declarations.

The type checking rule themselves are shown in Figure 19, 20 and 21. Also, see Clifton’s thesis [26] for details on these straightforward rules for standard OO expressions. Some rules we use the overriding union notation \uplus , defined in [70].

As in Clifton’s work [26, 23], the type checking rules are stated using a fixed class table (list of declarations) CT , which can be thought of as an implicit (hidden) inherited attribute. This class table is used implicitly by many of the auxiliary functions. For ease of presentation, we also follow Clifton in assuming that the names declared at the top level of a program are distinct and that the extends relation on classes is acyclic.

In the type checking rules above we use several auxiliary functions. Most of these are taken from Clifton’s dissertation [26, Figure 3.3]. A few others are given in Figure 22.

$\frac{}{\Pi \vdash n : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e == e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e != e' : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e > e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e < e' : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e >= e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e <= e' : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e + e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e - e' : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e * e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash !e : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e \&\&e' : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e' : \mathbf{exp\ int}}{\Pi \vdash e e' : \mathbf{exp\ int}}$	$\frac{}{\Pi \vdash e : \mathbf{exp\ c} \quad \text{isClass}(c)}{\Pi \vdash \text{isNull}(e) : \mathbf{exp\ int}}$	
$\frac{}{\Pi \vdash e : \mathbf{exp\ int} \quad \Pi \vdash e_2 : \mathbf{exp\ t_2} \quad \Pi \vdash e_3 : \mathbf{exp\ t_3} \quad t_2 \preceq t \quad t_3 \preceq t}}{\Pi \vdash \text{if}(e) \{e_2\} \text{else} \{e_3\} : \mathbf{exp\ t}}$		
$\frac{\text{isClass}(c) \quad c \not\preceq \text{Site}}{\Pi \vdash \text{new } c() : \mathbf{exp\ c}}$	$\frac{(\text{var} : \mathbf{var\ t}) \in \Pi}{\Pi \vdash \text{var} : \mathbf{exp\ t}}$	$\frac{\text{isClass}(c)}{\Pi \vdash \text{null} : \mathbf{exp\ c}}$
$\frac{\Pi \vdash e : \mathbf{exp\ c} \quad (c_2, t\ m(t_1\ \text{var}_1, \dots, t_n\ \text{var}_n) \{e\}) = \text{mbody}(c, m) \quad c \preceq c_2 \quad (\forall i \in \{1..n\} :: \Pi \vdash e_i : \mathbf{exp\ t'_i}) \quad (\forall i \in \{1..n\} :: t'_i \preceq t_i)}{\Pi \vdash e.m(e_1, \dots, e_n) : \mathbf{exp\ t}}$		
$\frac{}{\Pi \vdash e : \mathbf{exp\ c} \quad \text{fieldsOf}(c)(f) = t}}{\Pi \vdash e.f : \mathbf{exp\ t}}$		
$\frac{}{\Pi \vdash e : \mathbf{exp\ c} \quad \text{fieldsOf}(c)(f) = t \quad \Pi \vdash e' : \mathbf{exp\ t'} \quad t' \preceq t}}{\Pi \vdash e.f = e' : \mathbf{exp\ t'}}$		
$\frac{}{\Pi \vdash e : \mathbf{exp\ t} \quad \Pi \vdash e' : \mathbf{exp\ t'}}{\Pi \vdash e; e' : \mathbf{exp\ t'}}$		
$\frac{\text{isType}(t) \quad \Pi \vdash e : \mathbf{exp\ t'} \quad t' \preceq t \quad \Pi' = \Pi \uplus \{\text{var} : \mathbf{var\ t}\} \quad \Pi' \vdash e'' : \mathbf{exp\ t''}}{\Pi \vdash t\ \text{var} = e'; e'' : \mathbf{exp\ t''}}$		
$\frac{\text{isClass}(c)}{\Pi \vdash \text{cast } c\ e : \mathbf{exp\ c}}$	$\frac{(\text{loc} : \mathbf{var\ t}) \in \Pi}{\Pi \vdash \text{loc} : \mathbf{exp\ t}}$	$\Pi \vdash \text{NullPointerException} : \mathbf{exp\ \perp}$
$\Pi \vdash \text{ClassCastException} : \mathbf{exp\ \perp}$		

Fig. 20. Type-checking rules for standard expressions.

$$\begin{array}{c}
\text{(CHECK SERVICE)} \\
\frac{(\forall i \in \{1..n\} :: \text{isType}(t_i)) \quad (\forall j \in \{1..m\} :: \vdash \text{meth}_j \text{OK in } w)}{\vdash \text{service } w \{t_1 f_1; \dots t_n f_n; \text{meth}_1 \dots \text{meth}_m\} : \text{OK}} \\
\\
\text{(CHECK WEB-METHOD)} \\
\frac{\Pi' = \{var_1 : \text{var int}, \dots, var_n : \text{var int}\} \quad \Pi' \uplus \{\text{this} : \text{var } w, \text{thisSite} : \text{var } w\} \vdash e : \text{exp int}}{\vdash \text{int } m(\text{int } var_1, \dots, \text{int } var_n)\{e\} : \text{OK in } w} \\
\\
\text{(WEB-METHOD CALL EXP TYPE)} \\
\frac{\Pi \vdash e : \text{exp } w \quad \text{isService}(w) \quad (w, \text{int } m(\text{int } var_1, \dots, \text{int } var_n)\{e\}) = \text{mbody}(w, m) \quad (\forall i \in \{1..n\} :: \Pi \vdash e_i : \text{exp int})}{\Pi \vdash m(e_1, \dots, e_n)@e : \text{exp } t} \\
\\
\begin{array}{cc}
\text{(REFINING EXP TYPE)} & \text{(UNDER EXP TYPE)} \\
\frac{\Pi \vdash e : \text{exp } t \quad \Pi \vdash \text{spec} : \text{specFor } t}{\Pi \vdash \text{refining spec } \{e\} : \text{exp } t} & \frac{\Pi \vdash e : \text{exp } t}{\Pi \vdash \text{under } e : \text{exp } t} \\
\end{array} \\
\\
\text{(EVALBODY EXP TYPE)} \\
\frac{\Pi \vdash e_1 : \text{exp } t \quad \Pi \uplus \{\text{result} : \text{var } t\} \vdash e_2 : \text{exp int}}{\Pi \vdash \text{evalbody } e_1 e_2 : \text{exp } t} \\
\\
\text{(EVALPOST EXP TYPE)} \\
\frac{\Pi \vdash e_1 : \text{exp } t \quad \Pi \uplus \{\text{result} : \text{var } t\} \vdash e_2 : \text{exp int}}{\Pi \vdash \text{evalpost } e_1 e_2 : \text{exp } t} \\
\\
\text{(SPEC EXP SPEC TYPE)} \\
\frac{\Pi \vdash e : \text{exp int} \quad \Pi \uplus \{\text{result} : \text{var } t\} \vdash e' : \text{exp int}}{\Pi \vdash \text{requires } e \text{ ensures } e' : \text{specFor } t}
\end{array}$$

Fig. 21. Type-checking rules for new Tisa features.

$$\begin{array}{l}
\text{isClass}(t) = (\text{class } t \dots) \in CT \\
\text{isService}(t) = (\text{class } t \text{ extends Site } \dots) \in CT \\
\text{isType}(t) = (t = \text{int}) \vee \text{isClass}(t)
\end{array}$$

Fig. 22. Auxiliary functions used in type rules.

The notation $\tau' \preceq \tau$ means τ' is a subtype of τ . It is the reflexive-transitive closure of the declared subclass relationships with the added fact that that \perp is a subtype of all class type expressions. The type \perp is used as the type of exceptions. This is formalized in Figure 23.

$$\begin{array}{cccc}
\text{(BASIS)} & & \text{(REF)} & \text{(TRANS)} & \text{(BOTTOM)} \\
\frac{(\text{class } c \text{ extends } d \{ \dots \}) \in CT}{c \preceq d} & & \tau \preceq \tau & \frac{\tau_1 \preceq \tau_2 \quad \tau_2 \preceq \tau_3}{\tau_1 \preceq \tau_3} & \frac{\text{isClass}(c)}{\perp \preceq c}
\end{array}$$

Fig. 23. Subtyping rules, adapted from [26, Figure 3.4].

A.3 Omitted Details on Soundness of Refinement

This section proves the soundness of refinement.

Lemma 3. *If the atomic propositions in a specification \mathcal{S} are $\mathcal{P}(\mathcal{S})$ and the atomic propositions in program P are $\mathcal{P}(P)$, and P refines the specification \mathcal{S} then $\mathcal{P}(\mathcal{S}) \subseteq \mathcal{P}(P)$.*

The proof of this lemma follows from construction of \mathcal{P} and structural refinement rules shown in Figure 12. The construction of \mathcal{P} picks all potential web-method calls as propositions and the refinement ensures that all web-method specifications in \mathcal{S} have a corresponding web-method declaration in P .

Lemma 4. *Let $P \in \text{program}$ be given. If t' is a path for P , then there are paths t'_{pre} and t'_{loop} such that $t' = t'_{pre} + t'_{loop}$, t'_{pre} has finite length, and each $(z', \delta') \in t'_{loop}$ occurs infinitely often in t'_{loop} .*

Proof Sketch: If t' has finite length, then let $t'_{pre} = t'$ and t'_{loop} be the empty path.

If t' has infinite length, then since P has only a finite number of expressions, it must loop at some point. Consider all the states that occur infinitely often in t' , and let t'_{loop} be the longest suffix of t' that contains only such states. Let t'_{pre} be the unique prefix of t' such that $t' = t'_{pre} + t'_{loop}$. ■

Lemma 5. *Let \mathcal{S} be a specification and let P be a program such that \mathcal{S} is refined by P . Let $t + (z_{n-1}, \delta_{n-1})$ be a path for \mathcal{S} and let $t' + (z'_{n-1}, \delta'_{n-1}) + (z'_n, \delta'_n)$ be a path for P . If $\delta_{n-1} \Rightarrow \delta'_{n-1}$, then there is some (z_n, δ_n) such that $t + (z_{n-1}, \delta_{n-1}) + (z_n, \delta_n)$ is a path for \mathcal{S} and $\delta_n \Rightarrow \delta'_n$.*

Proof Sketch: From the definition of path for P , we have that z'_{n-1} represents an expression in P and that there is a control flow relation from z'_{n-1} to z'_n . From the derivation rules for expressions in programs, we have the following cases.

Case **if-true**: z'_{n-1} represents the **if** expression and z'_n represents the true expression. Case **if-false**: z'_{n-1} represents the **if** expression and z'_n represents the false expression. Case **seq**: z'_{n-1} represents the first expression and z'_n represents the second expression in the sequence. Case **def**: z'_{n-1} represents the definition expression and z'_n represents the second expression in the variable definition. Case **refining**: z'_{n-1} represents the **refining** expression and z'_n represents the body expression of the **refining** expression. Case **web-method call**: z'_{n-1} represents the web-method call expression and z'_n represents the body expression of the web-method.

From the assumptions we have that \mathcal{S} is refined by P . From the refinement rules we have that for each expression represented by z'_{n-1} and z'_n above there is a corresponding expression se_{m-1} and se_m in \mathcal{S} and the structure of these expressions and their relative order is identical. By the construction of the FSM (Figure 11) we have that for each case above corresponding to se_{m-1} and se_m there is some state z_{m-1} and z_m in Z and $(z_{m-1}, z_m) \in R$. Thus $t_{m-1} = t'_{m-1} + (z_m, \delta_m)$ is a path for P . Also for all cases except web-method call, there are no new atomic propositions corresponding to z'_n and z_m in program and specification, thus $\delta_m \Rightarrow \delta'_n$ is vacuously true.

For the case web-method call by the construction of the FSM (Figure 11), we have that the new set of propositions $\delta_m = \{m@w\}$. From the refinement rule for web-method call, we have that identical web-method call occurs in the program. From the definition of a path for P , we have that for each such occurrence of a web-method call the new of propositions $\delta'_n = \{m@w\}$. Thus $\delta_m \Rightarrow \delta'_n$ holds. ■

Proof of 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: Suppose P refines S . Let t' be a path for P .

The proof is by transfinite induction, using the various cases discussed in Figure 12 that could generate t' . The well-ordering on paths that is used is that $t_1 < t_2$ if and only if t_1 is a finite, proper prefix of t_2 .

Base case: Let t' be the empty path. Then by definition of refinement, the empty path for S is refined by t' , so we can choose t as the empty path.

Inductive case: Let t' be a non-empty (and potentially infinite) path for P . We assume inductively that for all $t'_1 < t'$ there is some path t_1 for S such that $t_1 \sqsubseteq t'_1$. We must show that there is some path t for S such that $t \sqsubseteq t'$.

By Lemma 4, we can write $t' = t'_{pre} + t'_{loop}$, such that t'_{pre} has finite length, and each $(z', \delta') \in t'_{loop}$ occurs infinitely often in t'_{loop} . Let t'_{loop} be chosen so that it is the longest such path.

Now there are two cases, depending on whether t'_{loop} is empty.

If t'_{loop} is empty, then $t' = t'_{pre}$ and t' is finite. Since t' is non-empty, we can write $t' = t'_{n-1} + (z'_n, \delta'_n)$. Now there are two subcases.

The first subcase is if t'_{n-1} is empty. Then by the construction of the FSM (Figure 11), we know that the propositions that are assigned a truth value at the start of a path (i.e., δ_n) are top-level calls to web-methods. Suppose this is a call to a web method m . But by assumption the program's m refines the corresponding web method in S , hence there must be a z_n and δ_n such that $\delta_n \Rightarrow \delta'_n$, and so in this case $[(z_n, \delta_n)] \sqsubseteq [(z'_n, \delta'_n)] = t'$.

The second subcase is if t'_{n-1} is non-empty. By definition of $<$ for paths $t'_{n-1} < t'$. So from the inductive hypothesis we get a sequence t_{n-1} such that $t_{n-1} \sqsubseteq t'_{n-1}$. Since t'_{n-1} is finite, it has a last element $(z'_{n-1}, \delta'_{n-1})$ and $t'_{n-1} = t'_{n-2} + (z'_{n-1}, \delta'_{n-1})$. Since $t_{n-1} \sqsubseteq t'_{n-1}$ it must be that t_{n-1} is finite and non-empty. Hence there is some t_{n-2} such that $t_{n-1} = t_{n-2} + (z_{n-1}, \delta_{n-1})$. Since $t_{n-1} \sqsubseteq t'_{n-1}$ it must be that $\delta_{n-1} \Rightarrow \delta'_{n-1}$. Thus by Lemma 5, there is some (z_n, δ_n) such that $t_{n-2} + (z_{n-1}, \delta_{n-1}) + (z_n, \delta_n)$ is a path for S and $\delta_n \Rightarrow \delta'_n$. Letting $t = t_{n-1} + (z_n, \delta_n)$, we then have $t \sqsubseteq t'$. This ends the proof of the second subcase, when t'_{loop} is empty.

If t'_{loop} is non-empty, we can write it as $t'_{loop} = (z'_{n+1}, \delta'_{n+1}) + t'_{pp}$. Since t'_{pre} is also non-empty, we can write $t'_{pre} = t'_{n-1} + (z'_n, \delta'_n)$. By the inductive hypothesis we have that there is some t_{pre} such that $t_{pre} \sqsubseteq t'_{pre}$. As above we can write $t_{pre} = t_{n-1} + (z_n, \delta_n)$, and by the refinement relationship, we know that $\delta_n \Rightarrow \delta'_n$. Thus by applying Lemma 5 again, we have that there is some (z_{n+1}, δ_{n+1}) such that $t_{n-1} + (z_n, \delta_n) + (z_{n+1}, \delta_{n+1})$ is a path for S and $\delta_{n+1} \Rightarrow \delta'_{n+1}$. Thus $t_{n-1} + (z_n, \delta_n) + (z_{n+1}, \delta_{n+1}) \sqsubseteq t'_{n-1} + (z'_n, \delta'_n) + (z'_{n+1}, \delta'_{n+1})$.

Now t'_{loop} must be made up of some repetitions of a prefix t'_2 of t'_{loop} that starts with $(z'_{n+1}, \delta'_{n+1})$. This path t'_2 is also finite, and so we can find a subpath t_2 in S such that

$t_2 \sqsubseteq t'_2$, as above. We can then paste these together to produce a path t in \mathcal{S} such that $t \sqsubseteq t'$. ■

A.4 Omitted Details on Soundness of Policy Verification

The key idea in the proof of soundness for policy verification is to give a state exploration technique to verify that the policy is satisfied by the state machine constructed by the construction algorithm in Figure 11. Furthermore, we show that the output of our construction algorithm is a valid finite-state program.

Lemma 6. *Given a policy $\phi \in \Phi(\mathcal{S})$ one can build a Büchi automaton $\mathcal{B}(\neg\phi)$ such that the language accepted by that automaton $\mathcal{L}(\mathcal{B}(\neg\phi))$ is exactly the set of computations satisfying the formula $\neg\phi$.*

The proof of this Lemma automatically follows from the proof of Theorem 2.1 and 3.3. given by Vardi and Wolper [30, pp. 4,6].

Given a finite state program $(\mathcal{Z}, s_0, \mathcal{R}, \Delta)$ one can construct an equivalent Büchi automaton $(\sigma, \mathcal{Z}, s_0, \varrho, \Delta)$, where $\sigma = 2^{\mathcal{P}(\mathcal{S})}$, $z' \in \varrho(z, \delta)$ iff $(z, z') \in \mathcal{R}$ and $\delta = \Delta(z)$ [30, pp. 5].

Lemma 7. *Given two Büchi automata $(\sigma, \mathcal{Z}, s_0, \varrho, \Delta)$ and $\mathcal{B}(\neg\phi)$ one can construct an automaton that accepts $\mathcal{L}((\sigma, \mathcal{Z}, s_0, \varrho, \Delta)) \cap \mathcal{L}(\mathcal{B}(\neg\phi))$.*

The proof of this Lemma also automatically follows from Lemma 3.1 of [30], which in turn follows from [71].

From Lemma 6 and 7 it follows that given a finite state program $(\mathcal{Z}, s_0, \mathcal{R}, \Delta)$ and a policy $\phi \in \Phi(\mathcal{S})$, one can construct an automaton that accepts a language, which is empty when the finite state program satisfies the policy. This emptiness property is known to be solvable in linear-time [72].

Lemma 8. *For a specification \mathcal{S} , the production relation \rightsquigarrow of Figure 11 constructs a valid finite-state program.*

Proof Sketch: The key intuition behind the proof of this lemma is that from the hypothesis of the rules (IF), (SPEC), (DEF), and (SEQ) in Figure 11 one can see that each of these rules generates a finite number of states. Furthermore, each of these rules maintains the structure of the finite-state program. The rule (WEB METHOD CALL) is different as it can potentially allow recursion, and thus generate potentially infinite number of states. However, this is accounted for by the (WEB METHOD CALL FSM 1) and (WEB METHOD CALL FSM 2) rules, which check membership in the table NT passed into the rule. The (WEB METHOD CALL FSM 1) rule requires that there is not already a state associated with the web method being called in NT , and ensures that subsequent relations use a table (NT' in the rule) that has the particular method defined. If there is a definition in the table, the (WEB METHOD CALL FSM 2) rule, is used, which does not add a new state but simply reuses the one in the table. This makes sure that the state for a particular web method call is only added to the FSM once. ■

Proof of 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.

Proof Sketch: The proof follows from lemma 6, 7, and 8.