

# Checking the Correctness of Architectural Transformation Steps via Proof-Carrying Architectures

R. A. Riemenschneider

Computer Science Laboratory, SRI International, Menlo Park, CA, USA  
rar@csl.sri.com

**Key words:** Software architectures, architecture hierarchies, transformation, refinement verification, proof-carrying architectures

**Abstract:** The end product of architecting is an *architectural hierarchy*, a collection of architectural descriptions linked by mappings that interpret the more abstract descriptions in the more concrete descriptions. Formalized transformational approaches to architecture refinement and abstraction have been proposed. One argument in favor of formalization is that it can result in architectural implementations that are guaranteed to be correct, relative to the abstract descriptions. If these are correct with respect to one another, conclusions obtained by reasoning from an abstract architectural description will also apply to the implemented architecture. But this correctness guarantee is achieved by requiring that the implementer use *only* verified transformations, i.e., ones that have been proven to produce correct results when applied. This paper explores an approach that allows the implementer to use transformations that have not been proven to be generally correct, without voiding the correctness guarantee. *Checking* means determining that application of the transformation produces the desired result. It allows the use of transformations that have not been generally verified, even ones that are known to sometimes produce incorrect results, by showing that they work *in the particular case*.

## 1. INTRODUCTION

The process of specifying an architecture often begins by providing a very high-level description of it. This description characterizes the

---

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35563-4\\_35](https://doi.org/10.1007/978-0-387-35563-4_35)

architecture in terms of a few abstract components, perhaps the principal functions the system must perform and some data stores. These components are linked by abstract connectors, perhaps indicating dataflow or control flow relationships among the components. This abstract description provides an easily understood overview of the entire system architecture, but omits so much detail that it provides relatively little guidance to someone charged with implementing the architecture using programming-language-level and operating-system-level constructs. So the abstract description must be successively refined—with complex components and connectors decomposed into simpler parts, and abstract specifications of operations and relationships replaced by more concrete specifications—until an appropriate amount of detail has been added. It usually is desirable to continue the refinement until implementation-level constructs have replaced all the abstractions.

Alternatively, architecting a system can consist of assembling instances of reusable component and connector types selected from a library. Such libraries effectively make the implementation-level architecture more abstract, and reduce the conceptual gap between the requirements specification and the implemented architecture. Nevertheless, combining a large number of components and connectors in complex ways can easily result in an architecture that is hard to understand and analyze. So, it is desirable to generate more easily comprehensible abstract representations of the implementation-level architecture.

In either case, the end product of the architecting process is typically a collection of architectural descriptions, at different levels of abstraction and often in different styles (Garlan, Allen, & Ockerbloom 1994). The more abstract descriptions are linked to the more concrete descriptions by interpretation mappings. An interpretation mapping says how the abstractions are implemented.<sup>1</sup> It sends each sentence in the language of the abstract description to a corresponding sentence in the language of the concrete description. For example, the fact that some component  $a$  is implemented by components  $a_1, a_2, \dots, a_n$  would be indicated by mapping the sentence

Component( $a$ )

to the sentence

Component( $a_1$ )  $\wedge$  Component ( $a_2$ )  $\wedge$  ...  $\wedge$  Component( $a_n$ )

<sup>1</sup>For more details on characterizing implementation steps using interpretation mappings, see our earlier paper (Moriconi, Qian, & Riemenschneider 1995).

The collection of architectural descriptions and interpretation mappings that comprise the complete architectural specification is called an *architecture hierarchy*.

There are many advantages to formalizing refinement and abstraction in system development: a library of refinement or abstraction transformations provides a “corporate knowledge base” of standard, or preferred, development patterns; mechanizing the application of these transformations lessens the likelihood of clerical errors during the development process; reuse of the transformations will result in greater validation of the patterns they codify; and so on. But one of the most fundamental advantages of formalization is that it allows the average developer to produce abstraction hierarchies that are guaranteed to be consistent. In other words, the use of verified transformations in the development process will guarantee that abstractions accurately characterize implementations, albeit more abstractly. A verified refinement transformation is one that has been proven to produce a correct implementation of whatever it is applied to. A verified abstraction transformation is one that has been proven to produce a correct abstraction of whatever it is applied to.

Even if attention is restricted to the case of architectures, there is some debate as to exactly what *correct* should mean. We have proposed a somewhat stricter-than-usual criterion for correctness (Moriconi, Qian & Riemenschneider 1995), while others have argued that the standard criterion is preferable (Philipps & Rumpe 1997). For present purposes, any reasonable criterion that characterizes correctness in terms of preservation of truth will do perfectly well. The standard correctness criterion is that every consequence of the abstract description must be a consequence of the concrete description as well. More precisely, for every sentence  $A$  in the language of the abstract description, where  $T_1$  is the logical theory that formalizes the abstract description,

$$T_1 \vdash A \Rightarrow T_2 \vdash \mu(A)$$

where  $T_2$  is the theory that formalizes the concrete description, and  $\mu$  is the interpretation mapping that links the two theories.<sup>2</sup> A mapping  $\mu$  that satisfies this condition is called an *interpretation of  $T_1$  in  $T_2$* . Our proposed stronger criterion for purely structural descriptions replaces the conditional with a biconditional, i.e., requires that the interpretation mapping be a *faithful theory interpretation*. One might also employ weaker-than-standard

<sup>2</sup>Our earlier paper explains how to formalize structural descriptions of architectures as logical theories. Since structural descriptions are largely declarative, the process is quite straightforward.

criteria, where only some consequences of the theory—properties of special interest—need be preserved.

What all these criteria have in common is that they justify the use of formal reasoning about the architecture based on the more abstract descriptions. If some sentence is shown to be a formal consequence of the abstract architectural theory, the concrete theory is known to correctly implement the abstract theory, and the sentence is among those that the correctness criterion guarantees are preserved by the implementation, then the sentence is known to be a consequence of the concrete theory as well. It is correctness guarantees that link the results of abstract analyses to the real world.

The usual approach to producing a correctness guarantee is restricting the architect to the use of verified transformations. This approach suffers from a problem, in practice. Even given a fairly mature library of verified transformations, it would hardly be surprising if an architect found himself unable to perform a certain refinement or abstraction step that he believed to be correct because the required transformation has not been included in the library. Expecting the typical system architect to produce a formal proof that the step is correct is unrealistic, yet the presence of a single unverified implementation step in the hierarchy voids the correctness guarantee provided by the restriction to verified transformations. Is there any way to allow the user to include such arbitrary steps in the development of the architecture hierarchy, while maintaining a correctness guarantee?

## 2. PROOF-CARRYING ARCHITECTURES

Our solution to this problem is based on the notion of checking the correctness of steps in architecture hierarchy development. By *checking*, we mean automatically performing some calculation that shows the step is correct. Checking can be substantially simpler than verification, because it is focussed on a particular step. Verifying a transformation means showing that it *always* produces correct results, while checking a transformation step means showing that a correct result was obtained *in one specific case*. Thus, checking entirely avoids the sometimes difficult problem of characterizing the preconditions required for the transformation to produce correct results (Riemenschneider 1998).

Our initial approach to checking transformation steps was inspired by work on compilers that generate proof-carrying code (PCC) (Necula & Lee 1998). The basic idea is that, rather than attempting to prove the transformations performed by a compiler always produce code with certain desired properties, to generate a purported formal proof that the compiled

code has those properties as part of the code generation process. The purported proof can then be checked and, if it turns out to be a correct proof, it follows that the generated code has the desired properties. Thus, the emphasis is shifted from showing that compiler transformations are correct in general to checking that they produced correct results in individual cases.

The application of this idea to architectural transformation is straightforward. At some abstract level, the architectural description is proven to guarantee that the architecture has some desirable property,  $C$ . The interpretation mapping  $\mu$  that sends abstract level sentences to their implementations can also be applied to the proof of  $C$ . If the image of the proof under the implementation mapping turns out to be a correct proof that the implementation has  $\mu(C)$ , then, of course, the implementation has  $\mu(C)$ . Checking the transformed proof can, therefore, provide the desired correctness guarantee.

### **3. AN EXAMPLE: SECURE DISTRIBUTED TRANSACTION PROCESSING**

The idea of proof-carrying architectures can be illustrated by an example, based on our development of software architectures for secure distributed transaction processing (SDTP) (Moriconi, Qian, Riemenschneider & Gong 1997). These architectures extend X/Open's standard DTP architecture (X/Open Company 1993) by enforcing a simple "no read up, no write down" security policy. The primary result of our development efforts is a hierarchy that links an extremely abstract architectural description, shown in Figure 1, to three implementation-level descriptions written in a style that can be directly translated into a programming language such as Java using standard network programming constructs. The gap between the abstract SDTP architecture and each concrete SDTP architectures is filled by roughly two dozen descriptions—the exact number varies among the implementations—at intermediate levels of abstraction, linked in a chain by interpretation mappings.

We are in the process of formally proving the implementation-level architectures are secure by proving that the abstract description is secure, and proving that every interpretation mapping preserves security. One of the techniques that is being employed is showing that the interpretations incrementally transform the abstract-level security proof into implementation-level security proofs. The example below shows how the interpretation mapping associated with the first refinement step in all three chains transforms the abstract security proof into a slightly more concrete security proof.

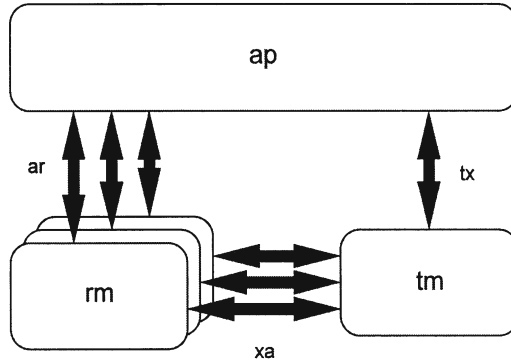


Figure 1. Abstract SDTP architecture — components linked by secure channels

### 3.1 The abstract SDTP architecture

Figure 1 depicts the most abstract architecture for SDTP. The boxes are the components of the architecture: the Application (labeled “ap”), some number of Resource Managers (labeled “rm”), and a Transaction Manager (labeled “tm”). The components are linked by *secure channels*, indicated by the heavy double headed arrows that make up the interfaces between the Application and Resource Managers, the Application and the Transaction Manager, and the Resource Managers and the Transaction Manager. Secure channels are a type of connector that enforce the security policy. In other words, secure channels will not carry classified data from a component to a component that lacks required clearances. To say that the system as a whole satisfies the security policy means that there is no flow of classified data to a component that lacks the required clearances.

### 3.2 An abstract proof of security

Informally, the security of the system follows almost immediately from the fact that it employs only secure channels. Not surprisingly, a textbook-style natural deduction proof (Lemmon 1987, Mates 1972) of system security is quite simple<sup>3</sup>. Consider the dataflow from some given Resource Manager *rm* to the Application *ap*, for example. A proof of the formula

<sup>3</sup> In this paper, I will use natural deduction, since that provides a familiar concrete representation of formal proofs. In our actual verifications that the SDTP hierarchy’s interpretation steps preserve security, we are employing the PVS verification system [18].

$$(\forall d : \text{Labeled\_Data}) [\text{Flows}(d, rm, ap) \supset \text{label}(d) \geq \text{clearance}(ap)]$$

which says

every labelled datum  $d$  that flows from  $rm$  to  $ap$  has a security label classifying it that is less than or equal to the clearance level of  $ap$

from five axioms of the architectural theory is shown in Figure 2.

{1}	1.	$(\forall d : \text{Labeled\_Data}) [\text{Flows}(d, rm, ap) \supset \text{Carries}(\text{secure\_ar\_channel}, d, rm's\_ar\_port, ap's\_ar\_ports(rm))]$	Axiom describing specific architecture
{2}	2.	$\text{Port\_Of}(ap's\_ar\_ports(rm), ap)$	Axiom describing specific architecture
{3}	3.	$(\forall c : \text{Secure\_Channel}) (\forall d : \text{Labeled\_Data}) (\forall x : \text{Output\_Port}) (\forall y : \text{Input\_Port}) [\text{Carries}(c, d, x, y) \supset \text{label}(d) \leq \text{clearance}(y)]$	Axiom characterizing secure channels
{4}	4.	$(\forall a : \text{Component}) (\forall y : \text{Input\_Port}) [\text{Port\_Of}(y, a) \supset \text{clearance}(y) \leq \text{clearance}(a)]$	Axiom constraining port clearances
{5}	5.	$(\forall s_1, s_2, s_3 : \text{Security\_Label}) [s_1 \leq s_2 \wedge s_2 \leq s_3 \supset s_1 \leq s_3]$	Axiom specifying transitivity of security label ordering
{1}	6.	$\text{Flows}(d_0, rm, ap) \supset \text{Carries}(\text{secure\_ar\_channel}, d_0, rm's\_ar\_port, ap's\_ar\_ports(rm))$	Universal instantiation (1)
{3}	7.	$\text{Carries}(\text{secure\_ar\_channel}, d_0, rm's\_ar\_port, ap's\_ar\_ports(rm)) \supset \text{label}(d_0) \leq \text{clearance}(ap's\_ar\_ports(rm))$	Universal instantiation (3)
{1, 3}	8.	$\text{Flows}(d_0, rm, ap) \supset \text{label}(d_0) \leq \text{clearance}(ap's\_ar\_ports(rm))$	Tautological consequence (6,7)
{4}	9.	$\text{Port\_Of}(ap's\_ar\_ports(rm), ap) \supset \text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap)$	Universal instantiation (4)
{2, 4}	10.	$\text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap)$	Tautological consequence (2,4)
{5}	11.	$\text{label}(d_0) \leq \text{clearance}(ap's\_ar\_ports(rm)) \wedge \text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap) \supset \text{label}(d_0) \leq \text{clearance}(ap)$	Universal instantiation (5)
{2, 4, 5}	12.	$\text{label}(d_0) \leq \text{clearance}(ap's\_ar\_ports(rm)) \supset \text{label}(d_0) \leq \text{clearance}(ap)$	Tautological consequence (10,11)
{1, 2, 3, 4, 5}	13.	$\text{Flows}(d_0, rm, ap) \supset \text{label}(d_0) \leq \text{clearance}(ap)$	Tautological consequence (8,12)
{1, 2, 3, 4, 5}	14.	$(\forall d : \text{Labeled\_Data}) [\text{Flows}(d, rm, ap) \supset \text{label}(d) \leq \text{clearance}(ap)]$	Universal generalization (13)

Figure 2. Formal proof that dataflow from  $rm$  to  $ap$  satisfies the security policy

The five axioms say

1. every labelled datum  $d$  that flows from  $rm$  to  $ap$  is carried by `secure_ar_channel` from the output port `rm's_ar_port` to the input port of the port array `ap's_ar_ports` that is indexed by  $rm$ ,
2. the input port of the port array `ap's_ar_ports` that is indexed by  $rm$  is a port of  $ap$ ,
3. if secure channel  $c$  carries labelled datum  $d$  from output port  $x$  to input port  $y$ , then  $d$ 's security label is less than or equal to the clearance level of  $y$ ,

4. the clearance level of any port  $y$  of component  $a$  must be less than or equal to the clearance level of  $a$ , and
5. the ordering of security labels is transitive.

The first two axioms are facts about the particular architecture, the third axiom is the defining property of the secure channel subtype, the fourth and fifth axioms are general axioms of the security model.

### 3.3 A slightly more concrete SDTP architecture

The secure channels of abstract SDTP architecture can be implemented in terms of ordinary dataflow channels and additional components in a variety of ways, depending upon the security properties of the components (Moriconi, Qian, Riemenschneider & Gong 1997). The most interesting implementation is shown in Figure 3, where the light double headed arrows represent ordinary dataflow channels that do not enforce the security policy.




Figure 3. More concrete SDTP architecture—secure channels refined to ordinary channels, or ordinary channels plus security filter

This implementation is suited to the case where all of the resource managers are single-level, but not necessarily the same level. The security policy is enforced by a multi-level secure component that filters dataflow between the application and the resource managers: if passing a datum from a resource manager to the application would violate the security policy, the filter removes it from the stream.



The concrete architecture can be thought of as resulting from the abstract architecture by applying several transformations. For example, one transformation, the *Filter Introduction Transformation (FIT)*, replaces secure channels between components that are not multilevel secure by ordinary dataflow channels and a component that enforces the security policy.

### 3.4 A slightly more concrete proof of security

Now it must be shown that, like the abstract SDTP architecture, the more concrete SDTP architecture has the desired security property. The two conventional approaches to establishing this result are

1. to directly prove that the more concrete architecture is secure, in much the same way the abstract architecture was proven secure (perhaps using the abstract-level proof for heuristic guidance), and
2. to show that the Filter Introduction Transformation (FIT), and the other transformations that produce the more concrete architecture from the abstract architecture, always preserves the security properly.

The use of proof-carrying architectures provides a third alternative.

When transformation FIT is applied, it can be applied not only to the architectural description, but to the formal security proof of Figure 2 as well. The result of applying FIT to this proof is shown in Figure 4, where the implementation mapping  $\mu$  associated with this application is determined as follows. A complete account of how first-order interpretation mappings are defined, and basic facts about them, can be found in logic textbooks (Enderton 1972, Shoenfield 1967)<sup>3</sup>. For present purposes, it is enough to know that

1. for every term  $\mathbf{t}$  of the language of the abstract theory,  $\mu(\mathbf{t})$  is a (possibly complex) term of the language of the more concrete theory,
2. for every predicate  $\mathbf{F}$  of the language of the abstract theory,  $\mu(\mathbf{F})$  is a (possibly complex) predicate of the language of the more concrete theory,
3. for every formula  $\mathbf{A}$  of the language of the abstract theory,

$$\mu(\neg\mathbf{A}) = \neg\mu(\mathbf{A})$$

and similarly for the other connectors, and

4. for every formula  $\mathbf{A}$  of the language of the abstract theory, every variable  $\mathbf{x}$ , and every type predicate  $\mathbf{T}$  of the language of the abstract theory,

$$\mu((\forall\mathbf{x} : \mathbf{T}) \mathbf{A}) = \mu(\forall\mathbf{x} : \mathbf{T}) \mu(\mathbf{A})$$

where  $\mu(\forall\mathbf{x} : \mathbf{T})$  is a sequence of universal quantifiers, and similarly for the other quantifiers.

<sup>3</sup> Technically, we will make use of what are called n-dimensional interpretations (Hodges 1993, pp. 212.) But this is a reasonably straightforward generalization of the definition found in the cited textbooks.

The Carries predicate

$$\text{Carries}(\langle \text{secure channel} \rangle, \langle \text{datum} \rangle, \langle \text{out port} \rangle, \langle \text{in port} \rangle)$$

that is mentioned in formulas 1, 3, 6, and 7 of the abstract-level proof is mapped to a conjunction of the Carries, Passes, and Carries predicates,

$$\text{Carries}(\langle \text{channel} \rangle, \langle \text{datum} \rangle, \langle \text{out port} \rangle, \langle \text{filter in port} \rangle)$$

$$\wedge \text{Passes}(\langle \text{filter} \rangle, \langle \text{datum} \rangle, \langle \text{filter in port} \rangle, \langle \text{filter out port} \rangle)$$

$$\wedge \text{Passes}(\langle \text{channel} \rangle, \langle \text{datum} \rangle, \langle \text{filter out port} \rangle, \langle \text{in port} \rangle)$$

This clause in the definition of  $\mu$  says that a secure channel carrying a datum from some output port to some input port is implemented as a channel carrying the datum from the output port to some input port of a filter, passing the datum through the filter from the input port to some output port, and carrying the datum from output port of the filter to the input port<sup>5</sup>. This mapping is also applied to formula 3 in order to preserve the fact that formula 7 should follow from formula 3 by Universal Instantiation.

The universal quantifier over secure channels in formula 3,

$$(\forall \langle \text{secure channel variable} \rangle : \text{Secure\_Channel})$$

is mapped by  $\mu$  to universal quantifiers over channels and a universal quantifier over MLS components,

$$(\forall \langle \text{to-filter channel variable} \rangle : \text{Channel})$$

$$(\forall \langle \text{filter variable} \rangle : \text{MLS\_Component})$$

$$(\forall \langle \text{from-filter channel variable} \rangle : \text{Channel})$$

It is easy to check that the result of applying the FIT interpretation mapping  $\mu$  to the proof of security is a syntactically correct derivation of the

<sup>5</sup> This mapping would not be appropriate to apply to every occurrence of the Carries predicate in every derivation, because some secure channels in the abstract architecture may not be replaced by a combination of two channels and a filter in the concrete architecture. However, formulas 1, 6, and 7 of the proof specifically refer to what `secure_ar_channel` carries, and this secure channel *is* being implemented by two channels and a filter, so I will use this simpler interpretation for purposes of the example.

desired security property from formulas that are images of axioms of the more abstract architectural theory. Mapping  $\mu$  sends tautological consequence steps to correct tautological consequence steps, universal instantiation steps to correct universal instantiation steps, and universal generalization steps to correct universal generalization steps. So  $\mu$  has indeed mapped the formal abstract-level security proof to a concrete-level security proof, but not necessarily a proof *from axioms of the concrete architectural theory*.

- {1} 1.  $(\forall d : \text{Labeled\_Data})[\text{Flows}(d, rm, ap)$   
 $\supset \text{Carries}(rm\_to\_filter\_channel, d, rm's\_ar\_port, filter\_in\_port(rm))]$   
 $\wedge \text{Passes}(mls\_filter, d, filter\_in\_port(rm), filter\_out\_port(rm))]$   
 $\wedge \text{Carries}(filter\_to\_ap\_channel(rm), d, filter\_out\_port(rm), ap's\_ar\_ports(rm))]$
- {2} 2.  $\text{Port\_Of}(ap's\_ar\_ports(rm), ap)$
- {3} 3.  $(\forall c_1 : \text{Channel})(\forall f : \text{MLS\_Component})(\forall c_2 : \text{Channel})(\forall d : \text{Labeled\_Data})$   
 $(\forall x_1 : \text{Output\_Port})(\forall x_2 : \text{Output\_Port})(\forall y_1 : \text{Input\_Port})(\forall y_2 : \text{Input\_Port})$   
 $[\text{Carries}(c_1, d, x_1, y_1) \wedge \text{Passes}(f, d, y_1, x_2) \wedge \text{Carries}(c_2, d, x_2, y_2)$   
 $\supset \text{label}(d) \leq \text{clearance}(y_2)]$
- {4} 4.  $(\forall a : \text{Component})(\forall y : \text{Input\_Port})[\text{Port\_Of}(y, a) \supset \text{clearance}(y) \leq \text{clearance}(a)]$
- {5} 5.  $(\forall s_1, s_2, s_3 : \text{Security\_Label})[s_1 \leq s_2 \wedge s_2 \leq s_3 \supset s_1 \leq s_3]$
- {1} 6.  $\text{Flows}(do, rm, ap)$   
 $\supset \text{Carries}(rm\_to\_filter\_channel, do, rm's\_ar\_port, filter\_in\_port(rm))$   
 $\wedge \text{Passes}(mls\_filter, do, filter\_in\_port(rm), filter\_out\_port(rm))$   
 $\wedge \text{Carries}(filter\_to\_ap\_channel(rm), do, filter\_out\_port(rm), ap's\_ar\_ports(rm))$
- {3} 7.  $\text{Carries}(rm\_to\_filter\_channel, d, rm's\_ar\_port, filter\_in\_port(rm))$   
 $\wedge \text{Passes}(mls\_filter, do, filter\_in\_port(rm), filter\_out\_port(rm))$   
 $\wedge \text{Carries}(filter\_to\_ap\_channel(rm), do, filter\_out\_port(rm), ap's\_ar\_ports(rm))$   
 $\supset \text{label}(do) \leq \text{clearance}(ap's\_ar\_ports(rm))$
- {1, 3} 8.  $\text{Flows}(do, rm, ap) \supset \text{label}(do) \leq \text{clearance}(ap's\_ar\_ports(rm))$
- {4} 9.  $\text{Port\_Of}(ap's\_ar\_ports(rm), ap) \supset \text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap)$
- {2, 4} 10.  $\text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap)$
- {5} 11.  $\text{label}(do) \leq \text{clearance}(ap's\_ar\_ports(rm)) \wedge \text{clearance}(ap's\_ar\_ports(rm)) \leq \text{clearance}(ap)$   
 $\supset \text{label}(do) \leq \text{clearance}(ap)$
- {2, 4, 5} 12.  $\text{label}(do) \leq \text{clearance}(ap's\_ar\_ports(rm)) \supset \text{label}(do) \leq \text{clearance}(ap)$
- {1, 2, 3, 4, 5} 13.  $\text{Flows}(do, rm, ap) \supset \text{label}(do) \leq \text{clearance}(ap)$
- {1, 2, 3, 4, 5} 14.  $(\forall d : \text{Labeled\_Data})[\text{Flows}(d, rm, ap) \supset \text{label}(d) \leq \text{clearance}(ap)]$

Figure 4. Transformed formal proof that dataflow from  $rm$  to  $ap$  satisfies the security policy

### 3.5 Completing the proof

The image of the first axiom under  $\mu$  says that every labelled datum that flows from  $rm$  to  $ap$  is carried to the filter from  $rm$ , passed through the filter, and then carried to  $ap$  from the filter. Just as in the case of the first axiom, this is a fact about the particular architecture that is either an axiom of the concrete theory, or easily and automatically derivable from axioms of the concrete theory. The mapping  $\mu$  leaves the second axiom unchanged. This will certainly be an axiom of the concrete theory, as well as the abstract theory. The image of the third axiom is a bit more complex. It states that the combination of the two channels and the filter enforces the security property. It is quite unlikely that this would be among the chosen axioms of the concrete-level theory, since it is the filter alone, effectively, that is

enforcing security. Still, it is easy to see that this formula must be a consequence of axioms of the concrete theory: the security model requires that channels that do not enforce security can only connect ports with matching clearances, and one of the defining properties of an MLS component is that it only supplies data at an output port if the classification of the data is less than or equal to the clearance of the port. A formalization of this proof from particular axioms we use in the SDTP security verification is shown in Figure 5.

- {1} 1.  $(\forall c : \text{Channel})(\forall d : \text{Labeled\_Data})(\forall x : \text{Output\_Port})(\forall y : \text{Input\_Port})$   
 $[\text{Carries}(c, d, x, y) \supset \text{clearance}(x) = \text{clearance}(y)]$   
 Axiom specifying connection constraint imposed by security model
- {2} 2.  $(\forall f : \text{MLS\_Component})(\forall d : \text{Labeled\_Data})(\forall y : \text{Input\_Port})(\forall x : \text{Output\_Port})$   
 $[\text{Passes}(f, d, y, x) \supset \text{label}(d) \leq \text{clearance}(x)]$   
 Axiom characterizing MLS components
- {3} 3.  $(\forall x)(\forall y)(\forall z)[x = y \supset [z \leq x \equiv z \leq y]]$  Instance of identity axiom schema
- {1} 4.  $\text{Carries}(c_2, d_0, x_2, y_2) \supset \text{clearance}(x_2) = \text{clearance}(y_2)$  Universal instantiation (1)
- {2} 5.  $\text{Passes}(f_0, d_0, y_1, x_2) \supset \text{label}(d_0) \leq \text{clearance}(x_2)$  Universal instantiation (2)
- {1, 2} 6.  $\text{Carries}(c_1, d_0, x_1, y_1) \wedge \text{Passes}(f_0, d_0, y_1, x_2) \wedge \text{Carries}(c_2, d_0, x_2, y_2)$   
 $\supset \text{label}(d_0) \leq \text{clearance}(x_2) \wedge \text{clearance}(x_2) = \text{clearance}(y_2)$   
 Tautological consequence (3,4)
- {3} 7.  $\text{clearance}(x_2) = \text{clearance}(y_2) \supset [\text{label}(d_0) \leq \text{clearance}(x_2) \equiv \text{label}(d_0) \leq \text{clearance}(y_2)]$   
 Universal instantiation (3)
- {1, 2, 3} 8.  $\text{Carries}(c_1, d_0, x_1, y_1) \wedge \text{Passes}(f_0, d_0, y_1, x_2) \wedge \text{Carries}(c_2, d_0, x_2, y_2)$   
 $\supset \text{label}(d_0) \leq \text{clearance}(y_2)$   
 Tautological consequence (6,7)
- {1, 2, 3} 9.  $(\forall c_1 : \text{Channel})(\forall f : \text{MLS\_Component})(\forall c_2 : \text{Channel})(\forall d : \text{Labeled\_Data})$   
 $(\forall x_1 : \text{Output\_Port})(\forall x_2 : \text{Output\_Port})(\forall y_1 : \text{Input\_Port})(\forall y_2 : \text{Input\_Port})$   
 $[\text{Carries}(c_1, d, x_1, y_1) \wedge \text{Passes}(f, d, y_1, x_2) \wedge \text{Carries}(c_2, d, x_2, y_2)$   
 $\supset \text{label}(d) \leq \text{clearance}(y_2)]$   
 Universal generalization (8)

Figure 5. Proof of image of abstract-level formula 3 under  $\mu$  from axioms of concrete theory

Discovery of this proof is easy. The form of the desired conclusion—a conjunction of conditions on Carries and Passes in the antecedent, and the comparison of label to clearance in the consequent—immediately suggests the use of the axioms on lines 1 and 2 of the proof. So it should be quite plausible that the proof can be discovered without human intervention by the transformation system. The interpretation mapping  $U$  does not affect the images of the remaining two axioms; they remain general axioms in the security model. So, by combining the proof in Figure 5 with the proof in Figure 4, we obtain a proof of the security property from axioms of the concrete theory. Moreover, this proof is recognizably a formalization of our informal argument (Moriconi, Qian, Riemenschneider & Gong 1997, p. 890) that the concrete architecture satisfies the security policy.

### 4. GENERALIZING FROM THE EXAMPLE

The idea of using the architectural transformation to transform the proof that the more abstract architecture has a desired property into a proof that the more concrete architecture has the property worked well for this rather simple, but real-world, example. Is there any reason to believe that it will work equally well in other cases?

Recall that the standard criterion for correctness of an implementation mapping  $\mu$  of an abstract logical theory  $T_1$  in a more concrete theory  $T_2$  is that  $\mu$  must interpret  $T_1$  in  $T_2$ , i.e., it must be the case that, for every formula  $A$  in the language of  $T_1$ ,

$$T_1 \vdash A \Rightarrow T_2 \vdash \mu(A)$$

If  $\mu$  interprets  $T_1$  in  $T_2$ , an easy inductive argument shows that  $\mu$  maps formal proofs from  $T_1$  to formal proofs from  $\mu[T_1]$  that can be extended to proofs from  $T_2$ . If  $A$  is an axiom of  $T_1$ , then, since  $\mu$  is a theory interpretation,  $\mu(A)$  is derivable from  $T_2$ . Because  $\mu$  is defined so that connectives pass through it,  $\mu$  maps tautological consequence steps to tautological consequence steps. Similarly,  $\mu$  maps universal instantiation and universal generalization steps to universal instantiation and generalization steps, respectively. Thus,  $\mu$  maps formal proofs from abstract axioms to formal proofs from images of abstract axioms, and images of abstract axioms can always be proved from concrete axioms, as shown in Figure 6.

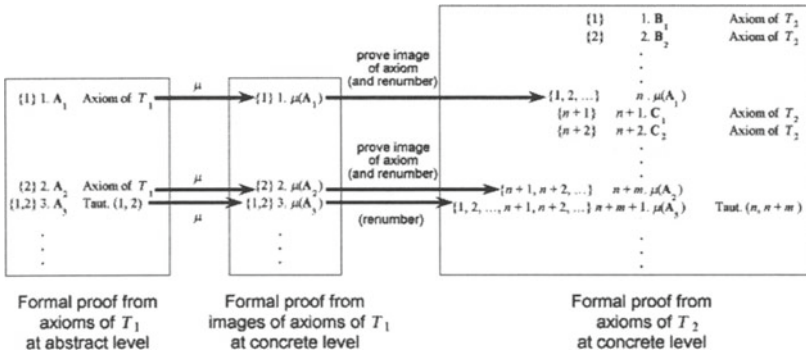


Figure 6. Interpretation of formal proofs

So, if an architectural transformation step is correct, in the standard sense, the corresponding interpretation mapping will map formal proofs to formal proofs containing gaps that can be filled. A fortiori, an abstract-level

formal proof of some particular property of interest—say, satisfaction of a security policy—will be mapped to a proof that the implementation also has (the implementation-level analogue of) the property. Since the replacement of the secure channel from *rm* to *ap* by a pair of channels and a filter is evidently correct, it is not surprising that the FIT mapping sends the abstract-level security proof to a concrete-level security proof.

It follows that the proof-carrying architecture approach allows the architect to perform arbitrary correct transformations when implementing an abstract architecture, provided the transformation system that supports the approach is clever enough to find the proofs of images of axioms.

The question remains: In general, how hard is it to discover these proofs? In our experience, it is invariably quite easy, because we deal with refinement patterns that make only small changes in representation of the architecture. Indeed, the example in Figure 5 is representative of the complexity of most of these proofs. At lower levels in the SDTP hierarchy, there are more gaps to be filled in—because lower-level architectural theories are more complex, and proofs are based on a larger number of axioms—but the size of the gaps is about the same. We are confident that considerable automated support for finding proofs to fill the gaps can be provided.

Finally, it should be noted that incorrect transformations that happen to preserve the proof of the property of interest will also be judged acceptable on the proof-carrying architectures approach. Therefore, it is well-suited to the case where the focus is on obtaining an implementation with some particular desirable property — i.e., when a weaker-than-usual correctness criterion is adequate — and placing minimal constraints on the architect's implementation options is preferred, as is the case in SDTP.

## 5. RELATED WORK

Although there is a large and growing literature on formal software transformation, nearly all of it is oriented toward maintaining functional correctness, rather than system structure. Similarly, there is a large body of literature on architectural refinement and composition, nearly all of it employing semiformal representation and analysis techniques, at best. The comparatively few papers on formal refinement of architectural structure include Broy's work on component refinement (Broy 1992), Brinksma, et al.'s, work on connector refinement (Brinksma, Jonsson & Orava 1991), Philipps and Rumpe's recent work on refinement of information flow architectures (Philipps & Rumpe 1997), and the work described in our own earlier papers. Also closely related is work by Garlan's group (Abowd,

Allen, and Garlan 1995), Luckham's group (Luckham, Augustin, Kenney, Vera, Bryan & Mann 1995) , and Moriconi and Qian's work on formally representing the semantics of connectors and relating semantic models at different levels of abstraction (Moriconi & Qian 1994). But, the emphasis in all these cases has always been on verification of general refinement patterns, rather than checking particular steps.

Necula and Lee's work on proof-carrying code and its applications (Necula & Lee 1996, 1997, 1998) introduced the notion of replacing verification by checking in the context of compilation. The work described in this paper can be viewed as generalizing their ideas about code refinement transformations to architectural transformations, both refinements and abstractions.

## 6. CONCLUSIONS

Transformational development of architectures can guarantee that implementations are correct by restricting the architect to a stock of verified transformations. But such a correctness guarantee is quite brittle, since use of a single non-verified transformation voids it. Moreover, if many transformations are used, and the verification of each is difficult, then confidence in the correctness of the implementation may be less than desired. Checking particular refinement steps offers a way of allowing the architect greater freedom, and of achieving higher levels of confidence that the implemented architecture has the desired properties.

Our initial approach to checking, based on the idea of proof-carrying architectures, is especially well suited to the case where the main requirement is high confidence that the implementation has some specific property. The property is shown to hold at some abstract level, and every refinement is produced by application of a transformation known to preserve the property, or is checked for correctness by making sure that the transformation preserves the proof of the desired property, or both.

The main limitation of this first approach to checking is that properties are checked one at a time. We are exploring other approaches to checking that allow an entire class of properties to be checked at once. One that seems particularly promising is based on the idea of applying the simplified technique for proving implementation mapping correctness (Riemenschneider 1997) to development steps at architecture definition-time. This complementary approach to checking will allow the correctness of steps to be checked, relative to our strong correctness criterion, rather than checking one or a few properties of interest. But it can be applied only to complete architectural descriptions of single structures, not to descriptions of



varied families of architectural structures. The proof-checking architectures approach applies equally well to descriptions of single structures and descriptions of families.

As mentioned above, our preliminary experiments with proof-carrying architecture are being performed with the PVS verification system (Owre, Rushby & Shankar 1992). Improved support for working with proof-carrying architectures, including automated discovery of the gap-filling proofs, is being implemented as part of the Xform<sup>4</sup> system, an enhanced version of our present architectural correctness checking toolset. *Xform*, pronounced *transform*, is a recursive acronym for “*Xform*, *for orderly reification*<sup>5</sup> and *maintenance*.” Xform will support transformational development and maintenance of architectural descriptions written in languages such as SADL (Moriconi & Riemenschneider 1997) and ACME (Garlan, Monroe & Wile 1997).

## ACKNOWLEDGEMENTS

This research was supported by the Defense Advanced Projects Research Agency (DARPA) Information Technology Office (ITO) under contracts F30602-95-C-0277 and F30602-97-C0040, whose support is hereby gratefully acknowledged. I would also like to thank Axel van Lamsweerde for his many helpful comments on the first draft of this paper.

## REFERENCES

- Abowd, G., Allen, R., and Garlan, D., (1995) Formalizing style to understand descriptions of software architecture. Tech. Rep. CMU-CS-95-111, School of Computer Science, Carnegie Mellon University.
- Brinksma, E., Jonsson, B., and Orava, F., (1991), Refining interfaces of communicating systems. *Proceedings of TAPSOFT '91*, S. Abramsky and T.S.E. Maibaum, Eds., Springer-Verlag, pp. 71—80
- Broy, M. (1992), Compositional refinement of interactive systems. Tech. Rep. No. 89, Digital Systems Research Center, Palo Alto.
- Enderton, H. B. (1972), *A Mathematical Introduction to Logic*. Academic Press.
- Garlan, D., Allen, R., and Ockerbloom, J. (1994), Exploiting style in architectural design environments. In *Proceedings 2nd ACM SIGSOFT Symposium on Foundations of Software Engineering SIGSOFT '94*, ACM Press, pp. 179—185.
- Garlan, D., Monroe, R.-T., and Wile, D. (1997), Acme: An architectural description interchange language. In *Proceedings of CASCON '97*. Available at

<sup>4</sup>*Xform* is a common mathematical shorthand for *transform*.

<sup>5</sup>*To reify* means to make actual. Thus, reification of software architectures is the process of turning them into actual implementations.



- <http://www.cs.cmu.edu/afs/cs/project/abel/www/acme-web/v3.0/white-paper-v3.0/white-paper.html>.
- Hodges, W. (1993), *Model Theory*. Cambridge University Press.
- Lemmon, E. J. (1987), *Beginning Logic*, second ed. Chapman and Hall.
- Luckham, D. C., Augustin, L. M., Kenney, J. J., Vera, J., Bryan, D., and Mann, W. (1995), Specification and analysis of system architecture using Rapide. *IEEE Transactions on Software Engineering* 21, 4, pp. 314—335.
- Mates, B. (1972), *Elementary Logic*, second ed. Oxford University Press.
- Moriconi, M., and Qian, X. (1994), Correctness and composition of software architectures. In *Proceedings 2nd ACM Symposium on Foundations of Software Engineering (SIGSOFT '94)*, ACM Press, pp. 164—174.
- Moriconi, M., Qian, X., and Riemenschneider, R. A. (1995), *Correct architecture refinement*. *IEEE Transactions on Software Engineering* 21, 4, 356—372. Available at <http://www.csl.sri.com/sadl/tse95.ps.gz>.
- Moriconi, M., Qian, X., Riemenschneider, R. A., and Gong, L. (1997), Secure software architectures. In *Proceedings of the 1997 IEEE Symposium on Security and Privacy*, pp. 84—93. Available at <http://www.csl.sri.com/sadl/sp97.ps.gz>.
- Moriconi, M., and Riemenschneider, R. A. (1997), Introduction to SADL 1.0: A language for specifying software architecture hierarchies. Tech. Rep. SRI-CSL-97-01, Computer Science Laboratory, SRI International. Available at <http://www.csl.sri.com/sadl/sadl-intro.ps.gz>.
- Necula, G. C., and Lee, P. (1998), The design and implementation of a certifying compiler. Submitted to PLDI '98. Available at <http://www.cs.cmu.edu/~necula/pldi98.ps.gz>.
- Necula, G. C., and Lee, P. (1996), Proof-carrying code. Tech. Rep. CMU-CS-96-165, School of Computer Science, Carnegie Mellon University. Available at <http://www.cs.cmu.edu/~necula/tr96-165.ps.gz>.
- Necula, G. C., and Lee, P. (1997) Efficient representation and validation of logical proofs. Tech. Rep. CMU-CS-97-172, School of Computer Science, Carnegie Mellon University. Available at <http://www.cs.cmu.edu/~necula/tr97-172.ps.gz>.
- Owe, S., Rushby, J. M., and Shankar, N. (1992), PVS: A prototype verification system. In *11th International Conference on Automated Deduction (CADE)* (Saratoga, NY.), D. Kapur, Ed., vol. 607 of *Lecture Notes in Artificial Intelligence*, Springer-Verlag, pp. 748—752.
- Philips, J., and Rumpe, B. (1997), Refinement of information flow architectures. In *Proceedings of the First IEEE International Conference of Formal Engineering Methods (ICFEM '97)*, pp. 203—212. Available at [http://www4.informatik.tu-muenchen.de/papers/icfem\\_rumpe\\_1997\\_Publication.html](http://www4.informatik.tu-muenchen.de/papers/icfem_rumpe_1997_Publication.html).
- Riemenschneider, R. A. (1997), A simplified method for establishing the correctness of architectural refinements. SRI CSL Dependable System Architecture Group, Working Paper DSA-97-02. Available at <http://www.csl.sri.com/sadl/simplified.ps.gz>.
- Riemenschneider, R. A. (1998), Correct transformation rules for incremental development of architecture hierarchies. SRI CSL Dependable System Architecture Group, Working Paper DSA-98-01. Available at <http://www.csl.sri.com/sadl/incremental.ps.gz>.
- Shoenfield, J. R. (1967), *Mathematical Logic*. Addison-Wesley.
- X/Open Company. (1993), *Distributed Transaction Processing: Reference Model*. Apex Plaza, Forbury Road, Reading, Berkshire RG1 1AX, U.K., November 1993.