

Error Reporting Logic

Ciera Jaspan* **Trisha Quan***
Jonathan Aldrich*

June 2008[†]
CMU-ISR-08-120

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

[†]Originally written April 2008

*School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

This work was supported in part by NSF grant CCF-0546550, DARPA contract HR00110710019, the Department of Defense, Army Research Office grant number DAAD19-02-1-0389 entitled “Perpetually Available and Secure Information Systems”, the Software Industry Center at CMU and its sponsors, especially the Alfred P. Sloan Foundation, and a fellowship from Los Alamos National Laboratory.

Keywords: first order logic, specification, error messages, error reporting

Abstract

When a system fails to meet its specification, it can be difficult to find the source of the error and determine how to fix it. In this paper, we introduce error reporting logic (ERL), an algorithm and tool that produces succinct explanations for why a target system violates a specification expressed in first order predicate logic. ERL analyzes the specification to determine which parts contributed to the failure, and it displays an error message specific to those parts. Additionally, ERL uses a heuristic to determine which object in the target system is responsible for the error. Results from a small user study suggest that the combination of a more focused error message and a responsible object for the error helps users to find the failure in the system more effectively. The study also yielded insights into how the users find and fix errors that may guide future research.


```

forall conn : ORMConnector |
  forall comp : Component |
    forall p : Port in comp.ports |
      (attached(conn.caller, p) ->
        declaresType(comp, Data) and
        declaresType(p, DataPort))

```

Figure 1: Sample Acme Specification

1 Introduction

Many specification languages are based upon first-order predicate logic. This is a very natural route to take for specifications; it provides a concise, expressive, and well-understood way for describing system-level details. Examples of such specification languages in recent literature include Acme, SCL, and Alloy [4, 8, 9]. In each of these languages, FOPL-based specifications constrain a system, and a tool produces errors when there is an inconsistency between the specifications and the system. The error messages produced by these systems generally fall into three categories:

- *Specification identifier.* Under this mechanism, the tool produces an error message that states which specification failed. The user must read the specification and manually analyze the system to determine which part of the system broke the specification.
- *Human generated message.* This mechanism attempts to provide the user with an intuitive understanding of the specification. The specification writer makes a generic summary about what the specification is checking, and this is used as the error message. The user can then use take message as a guide to understand the general problem.
- *Hybrid systems.* Some tools also hybridize the two mechanisms; they will use a human generated error message if it exists, but they will fall back on a specification identifier.

These mechanisms work very well for specifications that are short and have an obvious point of failure. However, they do not work well for complex specifications, such as the Acme specification shown in Figure 1. By Acme standards, this is a medium sized specification. It has 3 levels of quantification, a very small inner predicate, and it only calls pre-defined atomic predicates.

If the user must read the specification itself, they can quickly become lost in the details of the specification. There is no way to tell which sub-predicates in the specification failed, so the user must check each one. The user also doesn't know which objects in the system caused this failure.

Even if the specification writer provided an error message, this would not necessarily help a user. An error message would tell us the purpose of the specification, and this might help us look for bad patterns of behavior in our system. However, it still does not tell us which predicate failed or which object in our system caused the failure.

In Figure 1, the user would have to check the entire system for conformance to the specification. What we would really like is an error message that says:

*myPort must declare the type DataPort since
myConn.caller is attached to myPort*

Error reporting logic (ERL) provides an automated way for creating error messages such as the one above. ERL presents each failing point as a unique error. To do this, it singles out only the failing predicates and assigns responsibility of the error to a specific object in the system.

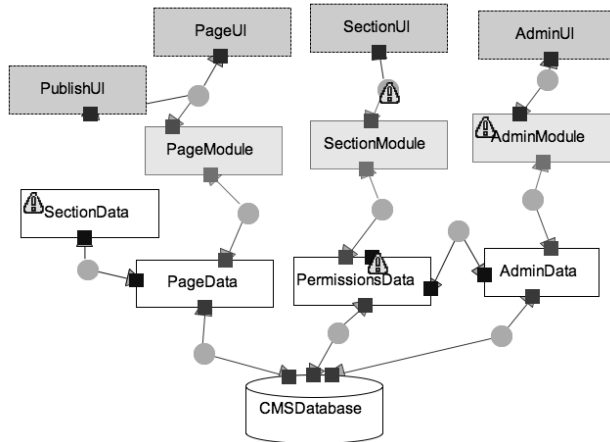


Figure 2: The Web System in AcmeStudio

In this paper, we will provide four contributions related to error messages from FOPL-based specifications:

- We present a user study that provides several insights into how users examine errors to find the root cause of the problem and how users attempt to fix the error. Primarily, we found that users see an error message as a single task which they must resolve, they only use keywords to find the problem rather than reading anything in depth, and they frequently rely on pattern recognition to find and fix errors. (Section 2)
- We present ERL, a system for automatically generating error messages from an existing specification based on first-order predicate logic. Section 3 will show how the ERL handles each of the specifications from the user study, and Section 6 shows how the implementation of ERL performs with MDS, the most complex architectural specification built with Acme.
- We have implemented ERL as a reusable component and have integrated it within AcmeStudio. The integration was relatively straightforward and required only a small amount of work to change the error messages. Section 4 provides implementation and integration details.
- During the user study, the same participants also used ERL. In section 5 we describe how the users reacted to the new error messages. Three of the four participants benefited from the ERL error messages. The remaining participant did not benefit, but was not hindered, by the error messages.

Throughout this paper, we will use the Acme specification language (and AcmeStudio, the graphical interface and checker for Acme) as our examples. AcmeStudio allows developers to view a graphical representation of an architecture. While the developers can access and edit the Acme code behind the graphical view, it is typically not used. AcmeStudio displays the architecture using component connector diagrams which can be edited entirely through a user interface. A sample diagram for an architecture is shown in Figure 2.

If an architecture fails to meet a specification, a red error triangle appears at the place where the specification was defined, as shown in Figure 2. This is not necessarily the component which is causing the failure. If the specification was defined at the system level, rather than the component level, then no error triangle appears.

Table 1: Participants

| ID | Configuration 1 | Configuration 2 |
|----|-----------------|-----------------|
| A | Web + ERL | Build |
| B | Web | Build + ERL |
| C | Build + ERL | Web |
| D | Build | Web + ERL |

In Acme, a software architect can choose to associate a handwritten error message to each specification. If the specification fails, for any reason, AcmeStudio displays the error message and a link to the specification code, in addition to the graphical indicator of the error. If the architect did not provide a default error message, then AcmeStudio displays only the link to the specification code. Since the architect can only write one message for the entire specification, the error message is typically about the general purpose of the specification.

2 How users find specification errors

We ran a small user study of AcmeStudio to determine how users fix errors, with and without ERL attached. The four participants had several months of classroom experience with AcmeStudio. The participants were told that the study was about usability in AcmeStudio and how developers find and fix errors; they were not told that the error messages were changed until after the main part of the study.

We provided the participants with two sets of Acme specifications, and we created an architecture for each which broke several of the specifications. The participants were asked to fix all the errors in the architectures.

Each participant used both sets of specifications, and each participant used AcmeStudio with and without ERL attached. Participants were each assigned a different configuration in a different order, as shown in Table 1.

The seeded errors were approximately equivalent in both systems. We created 5 categories of specifications, as detailed in Table 2. Each system contained a broken specification from each category, and they were approximately the same level of difficulty to find and fix.

While Acme supports hand written error messages, they are infrequently used in practice, and we did not include them in the user study. The participants either received a message from ERL, or they received the name of the specification which broke. In both systems, the specification source was easily available by double-clicking on the error. The specifications were written in a style familiar to Acme users, and we only used atomic predicates of Acme which our users were already familiar with.

- *binary relations* such as $<$, $>$, and $==$.
- *size(l)* to get the size of a list l .
- *attached($o1$, $o2$)* to test if $o1$ is directly attached to $o2$.
- *declaresType(o , t)* to test if o declares the type t .

Table 2: Broken specifications

| Type of error | How broken in the given architecture | Example specification from the user study |
|-----------------------------------------------------------------------------------------|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Simple</i> contains atomic predicates and at most 1 universal quantifier | The atomic fails once | rule atLeastOneAttachedRole = size(self .ATTACHEDROLES) >= 1; |
| <i>Conjunction</i> contains at least 1 conjunction of atomics and at most 2 quantifiers | Both parts of the conjunction fail in one instance | rule hasInAndOut = (exists p:Port in self .PORTS declaresType(p,ORMPort)) and (exists p:Port in self .PORTS declaresType(p,dataProvider)); |
| <i>Quantification</i> contains at least 2 quantifiers | The specification fails for two instances | rule ResultsOnly = forall comp:DeployResults in self .COMPONENTS forall p:inputPort in comp.PORTS declaresType(p, resultsPort); |
| <i>Disjunction</i> contains at least 1 disjunction and at most 1 quantifier | Both parts of the disjunction fail in one instance | rule usingXMLRoles = forall r:Role in self .ROLES declaresType(r, XMLReceiverRole) or declaresType(r, XMLProviderRole); |
| <i>Other</i> may contain any other predicate and one universal quantifiers ¹ | The predicate under test fails once. | rule compilingIsOutput = forall p:Port in self .PORTS declaresType(p, compilePort) -> declaresType(p, outputPort); |

2.1 Results from the control

We will start by looking at how the participants fixed errors in the control system. This provided several insights into problems users have with the existing error reporting mechanisms, including problems which ERL solves and problems for future work. In this section, we will look at trends we saw when participants from our user study used the original version of AcmeStudio. Later, in Section 5, we will see how ERL helped several users find the root cause of the error.

With each of the errors, all of the participants used the graphical cues as a starting point. Of the five specifications, only four had graphical cues; one rule was defined at the system level and therefore had no graphical cue. Participants B investigated these errors early on and decided to come back to them later because he could not easily find the location of the error.

The next step participants took was to read the specification source. When participants did this, they typically did not fully read the specification, but rather scanned it for keywords such as type names. When they flipped back to the architecture, they looked for a place where there were objects with those types all near each other. They would then investigate this area of the

¹BuildFamily had a failing implication, and WebFamily had a failing existential

architecture thoroughly to determine whether something was “obviously wrong”. The participants read the specification for the content it was checking for only when they could not find the problem through other means.

Participant D did attempt to read the specification to find an error, but soon ran into difficulties. The participant was working on the *conjunction* error from Table 2 and quickly read through the specification. The participant noted that there were two parts to this specification and said

“Does it tell me which side is failing?... Nope, no help.”

The participant then spent several minutes trying to understand the specification and reviewing the architecture around the error location. This participant became more frustrated at this point, asking

“But which part of the error is failing? It would be nice to see which part is failing so I don’t have to parse it.”

The final major technique participants used to find the error was pattern matching. They would start noticing the patterns of how architectural elements were laid out and then check other elements to see if they conformed to the same pattern. In some cases, participants would believe they had found a problem that didn’t actually exist, or they would find a problem that was different from the one they believed they were after. Upon finding any inconsistent patterns, the participant would attempt to make them identical.

This worked best if the participant understood the cause of the error before looking at the good example, or at least figured out the problem while they were looking at the example. If the participant did not understand the cause of the error, they could accidentally believe the correct example was actually the incorrect one. This problem occurred with the *quantification* errors from Table 2 since the participants could not tell which elements cause the failure.

Participant C used this technique quite frequently during the control part of the study. This participant made four comments about this during the study, usually comparing himself to a monkey:

“Doing like a monkey, trying to match patterns...”

In several cases, Participant C found the inconsistency and fixed the error without ever knowing what the problem was.

2.2 Expectations for ERL

Based on the information from our user study, we believe that error systems must:

- Direct users to the likely cause of the error, rather than the location where the specification is defined
- Assist users by including relevant keywords and excluding irrelevant ones
- Focus users on the part of the error they need to fix
- Provide users with examples that correctly pass the specifications

As we will see in the next section, ERL helps an error reporting system with 3 of the 4 objectives above. We leave the last task, providing users with examples, for future work, as it is not clear whether this will help users or possibly misdirect their attention.

$$\begin{aligned}
M.\text{text}(\Gamma, \text{declaresType}(a, b), \text{false}, \text{true}) &= \Gamma(a) + \text{“ must declare the type”} + \Gamma(b) \\
M.\text{text}(\Gamma, \text{attached}(a, b), \text{false}, \text{true}) &= \Gamma(a) + \text{“ must be attached to”} + \Gamma(b) \\
M.\text{text}(\Gamma, \text{equals}(a, b), \text{false}, \text{true}) &= \Gamma(a) + \text{“ must be equal to”} + \Gamma(b) \\
M.\text{text}(\Gamma, \text{equals}(a, b), \text{true}, \text{true}) &= \Gamma(a) + \text{“ must not be equal to”} + \Gamma(b) \\
M.\text{text}(\Gamma, \text{equals}(a, b), \text{false}, \text{false}) &= \Gamma(a) + \text{“ is equal to”} + \Gamma(b) \\
M.\text{text}(\Gamma, \text{equals}(a, b), \text{true}, \text{false}) &= \Gamma(a) + \text{“ is not equal to”} + \Gamma(b)
\end{aligned}$$

Figure 3: Sampling of atomic messages for Acme

3 Error Reporting Logic

In this section, we will use the error messages ERL produced for the study examples in Table 2 to explain how ERL breaks down specifications to include only the relevant information about an error. For each example, we will also look at the relevant ERL rule. The judgments for ERL are in the form

$$M, \Gamma \vdash p \hookrightarrow S$$

which is read as “Given the oracle M and context Γ , the predicate p produces the set of errors S ”.

Γ is a context that maps a unique variable name to a host-specific object. By *host*, we are referring to any FOPL-based specification system that uses ERL, such as Acme.

The *oracle* M is provided by the host specification system. The oracle provides answers to queries about *atomic predicates*, that is, a predicate which has some host-specific semantics. Our concept of an oracle is based on the concept of the oracle used in testing and [12]. In ERL, the oracle can be queried for the following:

- $evaluate(\Gamma, a)$ evaluates whether the atomic predicate a is true, given the context provided in Γ .
- $items(\Gamma, e)$ retrieves a list of objects for a quantifier, given some host-specific expression e .
- $text(\Gamma, p, isNegative, isDeontic)$ gets the message for the predicate p , give the context Γ . When p is an atomic predicate, this message is host-specific. We list a sampling of messages defined by the Acme oracle in Figure 3. If $isNegative$ is true, we must negate the message. If $isDeontic$ is true, the oracle produces a message in deontic mode (“a must be equal to b”), while if it is false, the message is stated as a fact (“a is equal to b”).

The set S is a set of pairs (r, m) where m is the error message and r is the *responsible object*, a host-specific object that ERL will blame the error on. Notice that for a single specification, the algorithm can produce multiple error messages, and each error message has its own responsible object. It is possible for the responsible object to have no value, represented in our rules as \bullet . In this situation, the host specification system may use its default assignment.²

The predicate p may be any first order logic predicate. ERL currently works for conjunction, disjunction, implication, negation, universal quantification, and existential quantification. Other first-order connectives, such as exclusive disjunction or unique quantification, can be added to

²Acme assigns the error to the object which defined the specification; this is the `self` object in the example specifications.

ERL, but higher order predicates are not supported. Predicates also include any atomic predicates that are defined by the host specification system. Atomics may be nested if the host system allows it, but ERL treats the entire predicate as an atomic and will not descend into it.

3.1 Simple Specifications

For the *simple* error shown in Table 2, ERL produces the error message:

The size of interData1.AttachedRoles must be greater than or equal to 1.

As the specification has only an atomic predicate, this was produced by directly querying the oracle for the truth of this statement and the error message. The message is stored in the error set. At the atomic level, we do not yet know which object will be “responsible” for this failure, so this is left as \bullet for now. Section 3.4 will show how this is filled in.

$$\frac{M.evaluate(\Gamma, a) = \text{true}}{M, \Gamma \vdash a \leftrightarrow \phi}$$

$$\frac{M.evaluate(\Gamma, a) = \text{false}}{M, \Gamma \vdash a \leftrightarrow \{(\bullet, M.text(\Gamma, a, \text{false}, \text{true}))\}}$$

3.2 Splitting errors

We would like to focus the user onto only the problems they need to fix and make each “fix task” independent. To do this, we will split errors upon evaluating a conjunction. By doing this, the *conjunction* error from Table 2 produced two error messages in the user study:

There must exist a p in SectionData.Ports such that p declares the type dataProvider.

There must exist a p in SectionData.Ports such that p declares the type ORMPort.

The ERL rule that produces these errors simply evaluates each side independently and produces two distinct errors. After the split, the messages may even have two different responsible objects. Each error represents a correction that the user must make in order to meet the specification.

$$\frac{M, \Gamma \vdash p_1 \leftrightarrow S_1 \quad M, \Gamma \vdash p_2 \leftrightarrow S_2}{M, \Gamma \vdash p_1 \wedge p_2 \leftrightarrow S_1 \cup S_2}$$

3.3 Joining errors

When the user attempts to fix the *disjunction* error in Table 2, they are working on a single task. Therefore, ERL shows a single error message. While the message can be lengthy, it contains all the keywords which a user might need to fix the error.

DataModelReceiver0 must declare the type XMLReceiverRole or DataModelReceiver0 must declare the type XMLProviderRole

This message was created by the ERL rule for joining messages on a disjunction failure. Notice that if we have already split the error on both sides, we must rejoin all the splits into a single error message. This does mean that the error messages are much longer, but they are also specific to

the task at hand. As an alternative to joining, if the specification system has hierarchical error reporting, ERL could create sub-errors and tell the user to fix one sub error.

$$\begin{array}{c}
\frac{M, \Gamma \vdash p_1 \hookrightarrow \phi \quad M, \Gamma \vdash p_2 \hookrightarrow S}{M, \Gamma \vdash p_1 \vee p_2 \hookrightarrow \phi} \\
\frac{M, \Gamma \vdash p_1 \hookrightarrow S \quad M, \Gamma \vdash p_2 \hookrightarrow \phi}{M, \Gamma \vdash p_1 \vee p_2 \hookrightarrow \phi} \\
\frac{M, \Gamma \vdash p_1 \hookrightarrow \{(r1_1, m1_1), \dots, (r1_k, m1_k)\} \quad M, \Gamma \vdash p_2 \hookrightarrow \{(r2_1, m2_1), \dots, (r2_j, m2_j)\} \quad k \geq 1 \quad j \geq 1}{M, \Gamma \vdash p_1 \vee p_2 \hookrightarrow \{(\bullet, m1_1 + \text{"and"} + \dots + \text{"and"} + m1_k + \text{" , or " } + m2_1 + \text{"and"} + \dots + \text{"and"} + m2_j)\}}
\end{array}$$

3.4 Assigning a responsible object

For a given failure, ERL uses a heuristic to determine which object in the system is at fault for the error. The heuristic states that the error should be assigned to the object bound in the nearest universal quantifier. Additionally, universal quantifiers split errors, so for the *quantifier* specification in Table 2, we get two messages:

InputT3 must declare the type resultPort.

(Responsible object is InputT3.)

InputT4 must declare the type resultPort.

(Responsible object is InputT4.)

$$\begin{array}{c}
\frac{M.items(\Gamma, L) = \phi}{M, \Gamma \vdash \forall x \in L . p \hookrightarrow \phi} \\
\frac{M.items(\Gamma, L) = \{o_1, \dots, o_n\} \quad M, \Gamma[x \mapsto o_1] \vdash p \hookrightarrow S_1 \dots M, \Gamma[x \mapsto o_n] \vdash p \hookrightarrow S_n \quad n \geq 1 \quad (x \text{ fresh in } \Gamma)}{M, \Gamma \vdash \forall x \in L . p \hookrightarrow \{(o_1, e) \mid (\bullet, e) \in S_1\} \cup \{(r, e) \mid (r, e) \in S_1 \wedge r \neq \bullet\} \\ \cup \dots \cup \\ \{(o_n, e) \mid (\bullet, e) \in S_n\} \cup \{(r, e) \mid (r, e) \in S_n \wedge r \neq \bullet\}}
\end{array}$$

Notice that if a failing predicate does not use the variable defined by the innermost quantifier, the responsible object may not appear in the error message (see Figure 4). A simple normalization of the specifications can fix this problem by examining both sides of a conjunction and pulling up a sub-predicate that does not reference the variable declared by the most recent quantifier.³ Figure 4 shows how a problem specification can be normalized to fix this problem.

If there was no universal quantifier surrounding the failing predicate, it is possible for ERL to return errors which do not have a responsible object. Since there are no universal quantifiers, the variables must either have been introduced by an existential or be pre-defined. In this case, ERL uses the default responsible object that would have been used by the host system. Acme defaults to the object which defined the specification, the `self` object, since this is the only known object.

³This problem does not occur with disjunction or implication because the sub-predicate that does not reference the variable will still appear in the same error message. Due to splitting, this is only a problem with conjunction.

Before normalization:

```
forall x in {A} |
forall y in {A,B} |
  exists z in {C,D} |
    x = y and passingAtomic(x,y,z)
```

After normalization:

```
forall x in {A} |
forall y in {A,B} |
  x = y and
  exists z in {C,D} |
    passingAtomic(x,y,z)
```

Error message: *A must be equal to B*

Responsible object without normalization: C

Responsible object with normalization: B

Figure 4: Normalization required

Universal quantifiers make a clear case for when splitting is important. Consider a specification which quantifies over a list of 100 elements, and 10 of these elements cause a failure, possibly failing in different ways. Instead of a single error, ERL will produce 10 errors. Each error would be associated to a distinct object, and the error messages themselves would be different forms if the specification failed in different ways for each variable binding.

3.5 Relying on current state

In most cases, ERL creates error messages in the deontic mode and describes a correction that the user must make to for the specification to be correct. However, ERL must sometimes describe the current state of the system to the user, such as in the rules for implication and existential quantification. In the rules for implication, ERL provides the user with information about how the error was triggered.

For the *other* error in Table 2, ERL produces:

outputT0 must declare the type outputPort since outputT0 declares the type compile-Port.

ERL produced this by asking the oracle for the text on the left side of the implication stated as a fact rather than as an instruction.

$$\frac{M, \Gamma \vdash p_1 \hookrightarrow \phi \quad M, \Gamma \vdash p_2 \hookrightarrow \phi}{M, \Gamma \vdash p_1 \implies p_2 \hookrightarrow \phi}$$

$$\frac{M, \Gamma \vdash p_1 \hookrightarrow S_1 \quad M, \Gamma \vdash p_2 \hookrightarrow S_2 \quad S_1 \neq \phi}{M, \Gamma \vdash p_1 \implies p_2 \hookrightarrow \phi}$$

$$\frac{M, \Gamma \vdash p_1 \hookrightarrow \phi \quad M, \Gamma \vdash p_2 \hookrightarrow S \quad S \neq \phi}{M, \Gamma \vdash p_1 \implies p_2 \hookrightarrow \{(r, e + \text{“since”} + M.\text{text}(\Gamma, p_1, \text{false}, \text{false})) \mid (r, e) \in S\}}$$

The rules for the existential quantifier also takes advantage of this message form. Like disjunction, exists must join the current error sets. While this results in a relatively longer message, it contains only the keywords that the user needs.

$$\begin{array}{c}
M.items(\Gamma, L) = \{o_1, \dots, o_n\} \\
M, \Gamma[x \mapsto o_1] \vdash p \hookrightarrow S_1 \dots M, \Gamma[x \mapsto o_n] \vdash p \hookrightarrow S_n \quad S_1 = \phi \vee \dots \vee S_n = \phi \\
n \geq 1 \quad (x \text{ fresh in } \Gamma) \\
\hline
M, \Gamma \vdash \exists x \in L . p \hookrightarrow \phi \\
\\
M.items(\Gamma, L) = \phi \quad (x \text{ fresh in } \Gamma) \\
\hline
M, \Gamma \vdash \exists x \in L . p \hookrightarrow \{(\bullet, \text{“There exists no ”} + x \\
+ \text{“ such that ”} + M.text(\Gamma, p, \text{false}, \text{false}))\} \\
\\
M.items(\Gamma, L) = \{o_1, \dots, o_n\} \\
M, \Gamma[x \mapsto o_1] \vdash p \hookrightarrow S_1 \dots M, \Gamma[x \mapsto o_n] \vdash p \hookrightarrow S_n \quad S_1 \neq \phi \wedge \dots \wedge S_n \neq \phi \\
n \geq 1 \quad (x \text{ fresh in } \Gamma) \\
\hline
M, \Gamma \vdash \exists x \in L . p \hookrightarrow \{(\bullet, \text{“There exists no ”} + x \\
+ \text{“ such that ”} + M.text(\Gamma, p, \text{false}, \text{false}))\}
\end{array}$$

3.6 Negation

ERL handles negation predicates separately from the other predicates. If simply we print out “not”, or an equivalent negative, anytime we see the predicate, we can introduce ambiguity and double (or more!) negatives. During normalization, ERL pushes not predicates inward to atomic predicates, and it requests that the oracle provide a reasonable negation messages for atomics. As most atomic messages are a single phrase, we have pushed the negatives down to a level where they are unambiguous and understandable.

$$\begin{array}{c}
M.evaluate(\Gamma, a) = \text{false} \\
\hline
M, \Gamma \vdash \neg a \hookrightarrow \phi \\
\\
M.evaluate(\Gamma, a) = \text{true} \\
\hline
M, \Gamma \vdash \neg a \hookrightarrow \{(\bullet, M.text(\Gamma, a, \text{true}, \text{true}))\}
\end{array}$$

4 Implementation of ERL

We implemented the ERL rules in Prolog, and we provided a Java wrapper and interface for the oracle. For a system to use ERL, it must be able to transform its specifications into the types defined by ERL, and it must provide an implementation of the oracle.

We implemented a transformer and oracle for AcmeStudio. The ERL addition to Acme required 139 LOC for the transforming functionality, and 643 LOC for the oracle. Of the lines of code for the oracle, 486 LOC were only for generating messages for atomic predicates and retrieving the names of elements in Γ . Acme utilizes all of the rules described in Section 3.

5 User study results

As discussed in Section 2, we ran a small user study where each participant attempted to fix errors in two Acme architectures. The users were provided with ERL for one of the two architectures. Both architectures contained 5 failing specifications, as described in Table 2, and ERL expanded these into 7 distinct errors due to splitting from conjunctions and universal quantifiers. The qualitative data suggests that ERL is helpful for many users, particularly for complex specifications. When it was not helpful, it did not misdirect or otherwise hinder users.

5.1 Results by type of error

For the *simple* errors, users did not receive any additional benefit from ERL. The graphical indicators were already in the correct place in the control configuration, and the specification was short enough that users could quickly find the problem. The errors were also fairly obvious from the rule names. Users almost always went directly to the cause and guessed what the problem was without reading the error message, so ERL did not help or hinder in this case.

One problem we noted was that fully qualified names in error messages confused participants. Upon seeing a qualified name, participants became “shell-shocked” by the number of words, so we have removed this from ERL. The graphical indicators already point the user to the location of the objects, so there should be little information lost. The participants in this study did receive error messages with fully qualified names, and we expect that this change would have improved the overall results.

As expected, ERL was much more helpful for conjunction errors. Participant D, who made several comments about not knowing which side of a conjunction was failing during the control portion of the study, was clearly helped by the ERL error messages. When using ERL, this same participant read the error message for the conjunction failures and fixed both errors in approximately three minutes.

The results of the disjunction error were surprisingly mixed. While we expected the wordiness to bother participants, participants A and D strongly preferred the ERL error message to using the system without ERL. When participant D initially opened the second system, he expressed concern that the errors were going to be as difficult to fix as before:

“Ugh, it’s all typecheck [errors] still...”

After examining a few error messages, the participant chose to start with the disjunction error and fixed it within a few minutes by doing what the error message suggested. Upon seeing the error go away, the participant commented:

“So, this seems like not too much thought.”

Participant C found that the error message was “not at all helpful”, though participant B did not find any of the error messages helpful.

ERL appeared to help participants A, C, and D when fixing errors that came from failing *existentials* and failing *implications*. In particular, ERL helped clear up confusion about variable bindings. In the control part of the study, Participant A was slightly confused by the “other” specification in Figure 2. The participant believed that two *different* ports had to be a compile port and an output port. Participant A read the specification and examined the seemingly correct system

several times before finally realizing the confusion. When participant C encountered this error with the ERL message instead, the participant did not even have the opportunity to be confused. The ERL error message replaced the variable `p` with the specific port name `outputT0`, and the participant clearly understood that this port had to be both an output port and a compile port. The only time participants saw variables in the error messages was when they encountered a failure from an existential. Participant D did not appear to be bothered by this, while participant A would jump to the source to understand the error better.

The true test of ERL was the *quantification* errors. These errors were generally the hardest to fix as they were the most complex, they were declared at the system level, and they failed in two places in the architecture. For these errors, ERL was clearly an improvement over the control system. In the control, participants narrowed down their search by reading the specification, but they still had problems after that. In the Build example, there were four objects that were being quantified over, and participants had to carefully inspect each one. They discovered the problem by carefully exploring each of the four objects and noticing that two were slightly different. Then they went back to the specification, determined which set of two objects were causing the problems, and corrected them. However, participant C believed that the correct connectors were the incorrect ones, and accidentally “fixed” the wrong connectors! The participant realized the mistake after the tool did not remove an error when rechecking the system. The ERL errors were clearly helpful in these cases, and participants appeared less frustrated during their search for the root cause.

5.2 Participant impressions of the ERL messages

After the users fixed the errors in both architectures, they took a post-survey about the error messages that they saw. Both participants C and D preferred the ERL messages. Participant A believed the two configurations were very similar and noticed little difference between the error messages. Interestingly, this participant used and was clearly helped by the error messages during the study. The error messages were possibly unobtrusive enough that the difference did not register to the user given all the other features of AcmeStudio.

Participant B preferred to just know which specification failed and view the source directly. However, the participant chose not to read the specification source when using ERL, even though the source was equally available in both systems. This participant switched quickly between tasks in both parts of the study and did not appear to spend much time focusing on the errors. The participant also completed very few tasks during the study and had to be stopped due to time constraints.

From this qualitative data, we believe that the error messages provided by ERL certainly help with some kinds of failures, and some users clearly prefer them. In *no* situations did ERL misinform the users, lead them away from the cause of the error, or otherwise hinder their progress. In each situation, it either helped or had no effect on their progress towards finding the error, other than a few seconds to read the message. For this reason, ERL has been put into use within AcmeStudio.


```

rule rule112 = R2_3(self)
  <<label : string = "Rule 2.3: An Actuator may only notify estimators of commands"
analysis R2_3(sys : system) : boolean =
  (forall compA : ActuatorT in sys.COMPONENTS |
  forall pA : CommandNotifReqrPortT in compA.PORTS |
  forall compX : Component in sys.COMPONENTS |
  forall pX : Port in compX.PORTS | connected(pA, pX) ->
    (declaresType(compX, EstimatorT) and declaresType(pX, CommandNotifProvPortT))

```

Figure 5: MDS Specification

6 Complex Examples from MDS

The Acme specifications in the case study were created for the purpose of the study, so in this section we explore how ERL handles a real Acme specification. For this purpose, we will use the Mission Data System (MDS), one of the most complex architectures specified in Acme. MDS specifies a state-based reactive control architecture for space systems. More about MDS and its Acme specification can be found in [3].

What makes MDS so complex is the number of constraints between two or more architectural elements. In order to express these constraints in Acme, the user needs a universal quantifier for each element, plus quantifiers over the sub-elements that attach larger elements together. The end result is that in order to specify a constraint between n elements, we may need $2n$ quantifications. For this reason, it is not uncommon for Acme specifications to have four or more quantifiers.

Given the complexity of these specifications, the writers of MDS also added generic error messages to each specification. In this section, we will compare these generic error messages to the specific error messages provided by ERL.

Figure 5 is a sample of a specification, and all necessary sub-specifications, from MDS. As we can see from the error message, this rule checks that only estimators receive commands from actuators. The specification `rule112` calls out to a sub-specification to do the work. If this specification fails because an Estimator's port was properly connected, but not of type `CommandNotifProvPortT`, the original version of AcmeStudio would give the error message:

Rule 2.3: An Actuator may only notify estimators of commands

Since the rule is defined at the system level, the user would have to investigate every connection between actuators and other components. To make matters more confusing, the user would probably look for an Actuator that is connected to something that it not an Estimator, when the real problem is the port type of the Estimator's port.

ERL would have produced the error message:⁴

estPort must declare the type CommandNotifProvPortT since actPort is connected to estPort.

and would direct the user to `estPort`, the port on the Estimator which is causing the failure. If this failure occurred multiple places in the system, then ERL would produce a distinct error for each failing port.

⁴The ability to descend into an Acme sub-specification is currently being implemented.

Another MDS rule checks that a component does not connect twice to a port on another component. Like the previous specification, this specification has several quantifiers and eventually has an implication that checks whether some ports are connected incorrectly. The generic error message for this rule is:

Rule 10: No two ports of a component should be connected to the same target port.

This does explain the problem that the specification is trying to find, but it doesn't tell us which ports are the problem. If a component had two ports, `portA` and `portB` that both eventually connect to `otherPort`, then ERL would produce the specific error message:

portA must equal portB since portA is connected to otherPort and portB is connected to otherPort.

Ideally, the user should see both the generic and specific messages. The generic description provides the user with the specification intent and would help the user understand the system goals. However, the ERL error message provides actionable guidance for how to fix the current error.

The last MDS example we consider checks that Sensors are in the correct state based upon how many Estimators are listening for data. The generic message for this rule is:

Rule 4.4: A sensor that it not connected to any estimators should specify that it is only raw data; if it is connected to more than one estimator, it should specify that it is informative to more than one.

Of course, only one of these two things could be true at any point; the sensor can not be hooked up to no estimators and more than one estimator at once. With ERL, not only does the user find out which sensor is causing the problem, they also find out which predicate is actually breaking and receive direct guidance on how to fix the error:

mySensor.rawData must equal true since the number of estimators connected to my-SensorPort is 0.

While the generic messages do help us understand the purpose of the specification and prevent us from making future errors, they do not help a user find the cause of their error. This is particularly important for specifications as complicated as MDS; even if the control system using MDS is small, it is still difficult to parse through the specifications by hand. If the system itself is also large, the user must spend a great deal of time re-checking parts of the system that are already correct. The user study and MDS examples show that ERL error messages are a useful addition to the existing error mechanisms because they help users to find the root cause of the error, even in complex specifications and large systems.

7 Related Work

Shapiro[12] explored and formalized algorithms for how programmers debug logic programs. Shapiro's algorithm for debugging a system with incorrect output is the most similar to the algorithm we have proposed. Like Shapiro, we investigate the sub-predicates for the source of the error, and we use an independent oracle to determine the correctness of a sub predicate. However, Shapiro's algorithm stops at the first failing sub predicate. Our algorithm continues to gather all of the failure points in the predicate, as well as produces them into a human readable error message.

Additionally, ERL uses a heuristic to identify a responsible object for the error so that the user receives direction on the failing object, not just the failing specification.

There is a large body of work on messages for typing errors (summarized in [7]). The research which uses program slicing [14] to find the causes of type errors [2, 6, 13] is the closest to ERL. Program slicing is a technique for analyzing which parts of a program are involved in computing the value of a variable at a particular program point. By analogy, ERL can be viewed as an approach for analyzing which parts of a specification and a model result are responsible for causing the specification to fail on that model. Rather than following data- and control-dependencies in a program, our approach analyzes how the truth of a logical specification depends on the truth of its parts.

Another system for describing typing errors, Seminal [11], uses a similar mechanism as ERL for separating the error-generation system from the checker itself. Seminal also treats the checker as an oracle of knowledge and will break down expressions into sub-expressions in order to find the root cause of a typing error. Upon finding the root cause, Seminal searches for similar sub-expressions that will typecheck, and it suggests the “best” similar sub-expression to the user as an alternative. However, the sub-expression produced by Seminal may not be the sub-expression the user actually wants, and may then mislead the user. While ERL does not currently provide a correct alternative, it also does not provide the user with misleading information. As the two systems provide different kinds of information, we expect that using both techniques would be beneficial for users.

ESC/Java uses an error reporting mechanism that also aims to provide the user with a failure point and a directed error message [10]. However, the error reporting mechanism is inherently different from ERL because ESC/Java’s checks that the specifications hold true universally as a set, while tools such as Acme check that individual specifications hold true. Since ESC/Java’s specifications must hold true together, the theorem prover can not break apart the specifications and check them individually the way ERL’s oracle does. It is the oracle’s ability to analyze sub-predicates of the specification that allow ERL to find the root cause of the error and provide the directed message. For ESC/Java, [10] can not find the root cause of the error, but it does display the point where the theorem prover found a counter-example to its proof. To show the user how it got into this bad state, the ESC/Java error generator creates a trace based upon labels it leaves in the logical predicates.

ESC/Java also has slightly different goals from ERL due to the way their users fix errors. The work on ESC/Java attempts to generate fewer errors and condense them; ESC/Java produce one error for each method rather than one error for each failing path. ERL attempts to do the opposite; it splits the errors at every opportunity. This difference makes sense when we consider how the users find and fix these errors. A user of ESC/Java works on the entire method and considers the whole problem one error. On the other hand, a user of Acme regards multiple failures from a universal quantifier as different errors. While the errors were all generated by the same specification, they are about different parts of the system and likely are not related.

The model checking community has also investigated error reporting [5]; the work which is closest to ours is that of [1]. The goals stated in [1] are very similar to the ones we present; they look to get at the cause of an error trace from a model checker, rather than the symptom. When

they determine the cause of the error, they then produce one error trace for each cause and generate separate error traces for each cause. The main hindrance is that, like the work with ESC/Java, it is difficult to treat a model checker as an oracle because it can not analyze sub predicates individually. Ball et al. proposes a heuristic for the problem by using correct traces to narrow down the problems in the failing traces. This heuristic does allow for a model checker to be treated in a fashion similar to our oracle, but it does require that enough correct traces exist to guide it.

8 Conclusion

We have presented error reporting logic (ERL), a system for automatically generating error messages from first-order predicate logic. ERL presents a user with a precise error message by automatically analyzing the specification to select only the predicates involved in the failure. Additionally, it uses a heuristic to assign fault to a particular object so that the user is directed to the point of failure.

Our user study shows that users were helped by the ERL error messages in certain cases, particularly in errors from conjunction, disjunction, and universal quantification predicates. ERL provided users with an indication of the source of the error and specific instructions about how to fix the error. When ERL did not help, it also did not mislead the users. This is a large improvement over the control system which did not provide the users with any specific guidance. While general guidance is useful for preventing future problems and providing knowledge for the user, it does not help the user fix the current problem.

The user study also provided some interesting insights into how users find the root cause of the error. In particular, we found that users frequently scan any text for keywords that will lead them to the cause of the error, and they only read text for content if they are stuck or want to confirm their suspicions. During the study, participants also vocalized concern about not knowing which parts of the specification was failing. Finally, we found that participants fell back to pattern matching when they could not be helped through other mechanisms. ERL addressed all the issues we saw except providing a “good” pattern to follow.

While ERL is certainly useful for Acme and similar specification systems, we anticipate that it will have greater benefit in more complex specification systems. Systems which require more complex logical connectives can easily extend the ERL concepts of splitting and joining errors to produce more useful error messages. ERL may also prove beneficial for systems where specifications are globally distributed by pinpointing only the relevant parts of the global specification. We look forward to seeing how other specification systems may be able to extend the concepts presented in ERL.

References

- [1] Thomas Ball, Mayur Naik, and Sriram K. Rajamani. From symptom to cause: localizing errors in counterexample traces. In *Principles of programming languages*, 2003.

- [2] V. Choppella and C.T. Haynes. Diagnosis of ill-typed programs. Technical Report 426, Indiana University, 1994.
- [3] D. Dvorak, R. Rasmussen, G. Reeves, and A. Sacks. Software architecture themes in JPL's Mission Data System. *IEEE Aerospace Conf. Proc.*, 7:259–268 vol.7, 2000.
- [4] David Garlan, Robert Monroe, and David Wile. Acme: an architecture description interchange language. In *Conf. of the Centre for Advanced Studies on Collaborative research*, 1997.
- [5] Alex Groce and Willem Visser. What went wrong: Explaining counterexamples. In *10th Intl. SPIN Workshop*, 2003.
- [6] Christian Haack and J. B. Wells. Type error slicing in implicitly typed higher-order languages. *Sci. Comput. Program.*, 50(1-3):189–224, 2004.
- [7] B.J. Heeren. *Top Quality Type Error Messages*. PhD thesis, Universiteit Utrecht, The Netherlands, 2005.
- [8] D. Hou and H.J. Hoover. Using SCL to specify and check design intent in source code. *Trans. on Software Eng.*, 2006.
- [9] Daniel Jackson. Alloy: a lightweight object modelling notation. *Trans. Softw. Eng. Methodol.*, 2002.
- [10] K. Rustan M. Leino, Todd Millstein, and James B. Saxe. Generating error traces from verification-condition counterexamples. *Sci. Comput. Program.*, 2005.
- [11] Benjamin S. Lerner, Matthew Flower, Dan Grossman, and Craig Chambers. Searching for type-error messages. In *Programming language design and implementation*, pages 425–434, 2007.
- [12] Ehud Y. Shapiro. Algorithmic program diagnosis. In *9th Principles of programming languages*, 1982.
- [13] F. Tip and T. B. Dinesh. A slicing-based approach for locating type errors. *Trans. Sfw. Eng. Methd.*, 10(1):5–55, 2001.
- [14] Mark Weiser. Program slicing. *Trans. Software Engineering*, July 1984.