

Protective Interface Specifications*

Gary T. Leavens¹ and Jeannette M. Wing²

¹ Department of Computer Science, Iowa State University, Ames, IA, USA

² Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA

Keywords: Specification languages; Underspecification; Partiality; Larch

Abstract. The interface specification of a procedure describes the procedure's behaviour using pre- and postconditions. These pre- and postconditions are written using various functions. If some of these functions are partial, or underspecified, then the procedure specification may not be well-defined.

We show how to write pre- and postcondition specifications that avoid such problems, by having the precondition “protect” the postcondition from the effects of partiality and underspecification. We formalize the notion of protection from partiality in the context of specification languages like VDM-SL and COLD-K. We also formalize the notion of protection from underspecification for the Larch family of specification languages, and for Larch show how one can prove that a procedure specification is protected from the effects of underspecification.

1. The Problem

This paper seeks to explain and precisely define properties of “good” procedure specifications. These properties say when the precondition of a procedure specification protects the postcondition from partiality or underspecification in the vocabulary used in the specification. While we will precisely define protection for formal specifications, it can be applied and used in even informal specifications (with, of course, less precision).

To explain what a protective specification is, we start with an informal

* An earlier version of this paper, without the appendix sections, appeared in Michel Bidoit and Max Dauchet (editors), *TAPSOFT '97: Theory and Practice of Software Development, 7th International Joint Conference CAAP/FASE, Lille, France*. Volume 1214 of Lecture Notes in Computer Science, Springer-Verlag, 1997, pages 520–534.

Correspondence and offprint requests to: Gary T. Leavens, Department of Computer Science, Iowa State University, Ames, IA 50011, USA.

```
int factorial(int x);
behavior {
  requires informally "x is not too big";
  ensures informally "result is the factorial of x";
}
```

Fig. 1. An ill-defined informal specification of a factorial procedure.

```
int factorial(int x);
behavior {
  requires informally "x is nonnegative and x is not too big";
  ensures informally "result is the factorial of x";
}
```

Fig. 2. A protective informal specification of a factorial procedure.

example. Consider an (ill-defined) specification of an integer-valued factorial procedure, such as that found in Fig. 1. This behavioural interface specification is to be implemented in C++, which explains why C++ syntax is used to specify how it is to be called. The pre- and postconditions follow **requires** and **ensures**, respectively; when the precondition is satisfied, the procedure must terminate in a state that satisfies the postcondition. (The keyword **informally** in Larch/C++ [Lea97] signals the start of an informal predicate.) This specification is ill-defined because it is not clear what the procedure should return when x is negative. The problem is that mathematics does not define what “the factorial of x ” means when x is negative, but for that case the specification seems to require a correct implementation to return some integer. Because of this, a correct implementation’s obligations when x is negative are unclear, making this a bad specification. Note that the problem with this specification has nothing at all to do with the particular mathematical formalism used to write the pre- and postconditions, or with any particular logic for reasoning about what they mean.

A better, yet still informal, specification of the factorial procedure is given in Fig. 2. In this specification the precondition requires that the argument x is nonnegative, and thus has a well-defined factorial. We say that the precondition of Fig. 2 “protects” the postcondition, because for all values of the arguments that satisfy the precondition, the vocabulary used in the post-condition is well-defined. Thus whatever the phrase “the factorial of x ” might mean when x is negative does not matter.

The concept of protection, even in informal specifications, does have one subtle twist. It is that one part of a precondition may protect other parts of the precondition itself, so that the entire precondition is well-defined. Most programmers are familiar with examples where they must check that a number is nonzero before checking some condition involving a ratio or modulo calculation. The same idea applies in specifications such as the one in Fig. 3, where the first conjunct in the precondition (“denom is positive”) protects the second. That is, if the first conjunct is false, the entire precondition is false, and so the meaning of the second conjunct does not matter, as the implementation will not have any specified behaviour in such a case.

In the example of Fig. 3, the (informal) logic used to reason about the meaning of the precondition matters. In our informal argument we assumed that if the first conjunct in the precondition is false, then the entire precondition is false (and hence well-defined). However, since the precondition is informal, one could plausibly argue that since the “/” operator used in the second conjunct is partial,

```
double taxFor(int base, int num, int denom);
behavior {
  requires informally "denom is positive and  $0 \leq (\text{num}/\text{denom}) \leq 1$ ";
  ensures informally "result is approximately  $(\text{num}/\text{denom}) * \text{base}$ ";
}
```

Fig. 3. A protective specification that demonstrates protection within the precondition.

it has no meaning when “denom” is zero, and in that case perhaps the entire precondition should be considered meaningless. To resolve such questions, one must take the first step towards a formal specification language, and agree on some conventions for interpreting such formulas.

In this paper we consider what protection means with respect to partiality and underspecification. Our treatment of protection is not meant to be exhaustive, but merely to illustrate concepts that are useful with some logics that are widely used for formal specification. (See [ChJ90, GrS95] for surveys that also cover additional kinds of logics that might be used in formal specification, and hence might need their own concepts of protection. Also PVS [ORS95] represents another kind of specification logic that should be considered in extending our concepts.)

The first concept of protection we discuss is appropriate for behavioural interface specification languages (BISLs) that use a logic that accepts the existence of partial functions and has various non-classical ways to reason about them. For example, VDM-SL [Jon90, A⁺93] uses a logic called LPF [Jon90, Section 3.3] [BCJ84, BFL94, JoM94], which has three logical values and two kinds of equality.¹ As another example, the specification language COLD-K [FeJ92] uses a logic having just two logical values, but in which all other types have an *improper* value, \perp , which models the “undefined” results of partial functions, and also models computations that go into infinite loops or cause errors. All other values are *proper*. In COLD-K there is also a definedness predicate, D , that allows one to reason explicitly about whether a term denotes a proper value or not. There are several other languages with similar concepts [Bij90, Bli91, SKT95, KTB91, WDC95].

The second concept of protection we discuss is appropriate for BISLs that use a logic that does not admit the existence of partial functions, but uses *underspecification*. In such a logic, one avoids specifying a value for undefined terms [GrS95, Jon95]. In this approach, to make a term “undefined” one simply does not specify its value; hence it will not be possible to prove anything about such a term. This kind of logic is used in the Larch Shared Language, LSL [GHG93, Chapter 4] [GHM90], which is the mathematical component of the Larch family BISLs [GHG93], in the BISLs of the RESOLVE family [OSW94], and in Z [Hay93, Spi92] (according to its draft standard [WoJ96]). The subtle problems that underspecification may cause for the unwary in LSL (and similar logics) are discussed in Appendix A; indeed Jones’s paper pointing out these problems [Jon95] motivated the present work.

It is not the purpose of this paper to advocate one kind of logic over another. Some logic that deals with partiality is needed to verify client code and implementations, because of the partiality that is common in pre- and postconditions. Instead, this paper explores concepts of protection, with the aim

¹ However, in LPF nonstrict (i.e., strong) equality and the definedness operator, Δ , are only used in meta-arguments, since the logic is designed so that one only needs to use strict (i.e., weak) equality in proofs.

```

fact: int -> int
fact(i) == if i = 0 then 1 else i * fact(i-1)

FACTORIAL(x: int) result: int
  pre 0 <= x and x <= 8
  post result = fact(x)

```

Fig. 4. An auxiliary function specification and a protective procedure specification for factorial in VDM-SL. (Note that the factorial of 9 is larger than 2^{16} .)

of improving intuition about it and providing more guidance to specifiers. We also discuss how to prove protection from the effects of underspecification.

2. Protective Procedure Specifications

The idea of protection in a BISL was first formulated by Wing [Win83, Section 5.1.4]. Although we generalize that notion here, our goal is the same as Wing's original: knowing when a behavioural interface specification protects "its users from the incompleteness of the" mathematical vocabulary used in that specification "by ensuring that the meaning of the procedure specification is independent of any incompleteness" in that vocabulary (p. 123). In addition, we ensure that the precondition protects itself.

2.1. Partiality Protection

In a specification language like VDM-SL, COLD-K, or TROLL *light* [GCD95], the notion of a procedure specification that protects against partiality is relatively straightforward. This is because the associated logic explicitly includes a "bottom" element, \perp , and a definedness predicate, which we will write as D (where $D(\perp) = \text{false}$ and if x is proper then $D(x) = \text{true}$). The symbol \vdash stands for provability in the appropriate logic (for LPF, at the metalogical level). The idea is that a specification is protective if for all possible inputs, the precondition is defined, and whenever the precondition is true, then the postcondition is defined. Note that the precondition has to protect itself because it must always be defined.

Definition 2.1 (partiality-protective) A procedure specification, S , that uses a mathematical theory, T , and has formal parameters, $\vec{x} : \vec{U}$, precondition, $Q(\vec{x})$, and postcondition, $R(\vec{x})$, is *partiality-protective* if and only if

- $T \vdash \forall \vec{x} : \vec{U} . D(Q(\vec{x}))$, and
- $T \vdash \forall \vec{x} : \vec{U} . Q(\vec{x}) \Rightarrow D(R(\vec{x}))$.

For example, the VDM-SL specification of factorial in Fig. 4 is partiality-protective, because the precondition is always defined, and whenever x satisfies the precondition, the postcondition is always defined.

2.2. Underspecification Protection

The Larch family, the RESOLVE family, and Z use logics in which all functions are total. Since we are most familiar with Larch, we concentrate on Larch in

```

bufferTrait: trait
  includes Integer
  introduces
    bufSize: → Int
  asserts
    equations
      0 < bufSize ∧ bufSize ≤ 1024;

```

Fig. 5. A trait with an underspecified constant.

the discussion below. The appropriate notions for RESOLVE and Z can be defined similarly. For a logic that regards all functions as total, the notion of partiality protection has no meaning. The analogous notion, which we call “underspec-protection,” is a test that the meaning of a procedure specification does not rely on underspecified terms. Note, however, that an operator may be underspecified for reasons other than being “partial.” For example, in Fig. 5, `bufSize` is underspecified but not partial in any sense.²

We define the notion of underspec-protection in three steps. First we define the notion of a primed LSL trait and primed LSL term. (An LSL *trait* is a specification of mathematical vocabulary in an augmented form of first-order logic with equality [GHG93, Chapter 4].) That notion is used to describe a notion of a “completely-defined” term. An LSL term is completely-defined if it can be proved to have the same value in all models of its trait. A completely-defined term is similar to a defined (non- \perp) term in logics like LPF; this is the main technical distinction between the two notions of protection. Finally we define the notion of underspec-protection itself.

The notion of a primed trait and term is a variation of the idea of “priming” traits and terms found in the Larch Prover (where it is used in proving that an operator is “converted” [GHG93, pp. 142–4]).

Definition 2.2 (Primed Trait, T') Let T be an LSL trait. Let T' be a version of the trait T with every operator f in T replaced by f' , except that the following operators are left alone:

- all operators in the built-in trait `Boolean`,
- all operators in all instances of the built-in traits `Conditional` (which specifies **if then else**), and `Equality` (which specifies the operators `=` and `≠`), and
- all operators mentioned in a **generated by** clause.

For example, consider the trait `factTrait`, given in Fig. 6. The trait `factTrait'` has `fact` replaced by `fact'`, but `true` and the boolean operators are not primed, and neither are `0`, `pred`, and `succ`, because they are mentioned in the **generated by** clause of the trait `Integer` [GHG93, p. 161]. Operators mentioned in a **generated by** clause are meant to give a way to produce all values of a given sort; priming these would add “junk” to the specification. Another reason for not priming operators mentioned in a **generated by** clause is that if one imagines constructing equivalence classes from terms, one starts with the terms formed from the generators, and collapses equivalent ones; in this process, a generator applied to the same arguments always ends up in the same equivalence class.

² In these logics, there is also no way to separate underspecification that is used to make operators “partial” from underspecification that is used to make specifications intentionally less constraining, as in a `choose` operator for sets.

```

factTrait: trait
  includes Integer
  introduces
    fact: Int → Int
  asserts
    ∀ i: Int
      fact(0) == 1;
      (i > 0) ⇒ (fact(i) = i * fact(i-1));

```

Fig. 6. A trait for factorial, written in LSL.

This reflects the model theory of LSL, in which operators are all (deterministic) functions. Hence it is not necessary for generators to be canonical; i.e., it is not necessary that the generation process be free. For the same reason it is fine to have more than one set of generators asserted for a type.

Similarly, if P is a term in the language of T , then let P' be a copy of P with every operator f that appears in P replaced by f' , with the same exceptions as for primed traits. For example, if P is “**result** = **fact**(x)”, then P' would be “**result** = **fact'**(x)” because **fact** is not exempted from priming, “=” is exempt from priming, and **result** and x are not operators. As another example, if P is “**bufSize**” from the trait in Fig. 5, then P' would be **bufsize'** because **bufSize** is an operator.

In what follows, we write $T \vdash P$ to mean that P is provable from trait T .

Definition 2.3 (completely-defined) An LSL term, $P(\vec{x})$, with free variables \vec{x} of sorts \vec{U} , is *completely-defined for trait T* if and only if

$$T \cup T' \vdash \forall \vec{x} : \vec{U}. P(\vec{x}) = P'(\vec{x})$$

Trivial examples of completely-defined terms include variables, because for each trait T , $T \cup T' \vdash \forall x : U. x = x$. A more interesting example is that, for **factTrait**, the term **fact**(27) is completely-defined, but both **fact**(-1) and **fact**(x), where $x : \text{Int}$, are not. As another example, consider the trait **ChoiceSet** [GHG93, p. 176] where the operator **choose** is specified as follows.

$$b \neq \{ \} \Rightarrow \text{choose}(b) \in b$$

For this trait, the term **choose**($\{1\} \cup \{2\}$) is not completely-defined.

The following definition of when a procedure specification is protective says, in essence, that the precondition must be completely-defined for the used trait, and hence must protect itself. In addition, whenever the precondition holds, then the postcondition must be completely-defined. The two requirements in the definition are analogous to those for partiality protection, with complete-definition tests playing the role of the definedness predicate.

Definition 2.4 (underspec-protective) A procedure specification, S , that uses trait T , has formal parameters $\vec{x} : \vec{U}$, precondition $Q(\vec{x})$, and postcondition $R(\vec{x})$, is *underspec-protective* if and only if

- $T \cup T' \vdash \forall \vec{x} : \vec{U}. Q(\vec{x}) = Q'(\vec{x})$, and
- $T \cup T' \vdash \forall \vec{x} : \vec{U}. Q(\vec{x}) \Rightarrow (R(\vec{x}) = R'(\vec{x}))$.

The definition of underspec-protective suggests a direct proof technique. For example, to prove that the specification of **factorial** in Fig. 7 is underspec-protective, one must show that **factTrait** \cup **factTrait'** proves both of the following:

```

uses factTrait(int for Int);

int factorial(int x);
behavior {
  requires 0 ≤ x ∧ x ≤ 8;
  ensures result = fact(x);
}

```

Fig. 7. A specification of the factorial procedure in Larch/C++.

- $\forall x : \text{int} . (0 \leq x \wedge x \leq 8) = (0 \leq' x \wedge x \leq' 8')$, and
- $\forall x : \text{int} . (0 \leq x \wedge x \leq 8) \Rightarrow (\mathbf{result} = \mathbf{fact}(x)) = (\mathbf{result} = \mathbf{fact}'(x))$

The logic used to do these proofs has to handle the “unspecified” values of $\mathbf{fact}(x)$ when x is negative; this is no problem in the logic of LSL and the other logics considered in this section.

Proofs, such as the one sketched above, that a procedure specification is underspec-protective are quite tedious to carry out in detail, at least by hand. While they may be amenable to machine support, it is also convenient to define a notion that is easier for humans to deal with.

3. Proving Underspec-Protection

In this section we describe an easier way to prove underspec-protection in a Larch family BISL. This proof technique uses extra information that specifiers could add to LSL traits. This extra information would also allow a user of LSL to specify more precisely and check what is intended to be completely-defined.

Since we are only concerned with underspec-protection in this section and the next, we will simply refer to it as “protection” in informal remarks.

3.1. Specifying What is not Underspecified

LSL already has some provision for specifying what is not underspecified — the specification of when an operator is “converted.” This is done by using a **converts** clause. A **converts** clause says that the axioms of the trait uniquely define the operators named in the clause, “relative to the other operators in the trait” [GHG93, p. 142]. (We include in Appendix B a more detailed explanation of conversion for the sake of completeness.)

However, proving that an LSL operator is converted does not mean it is completely-defined; it may still be underspecified. For example, consider the trait in Fig. 8. In this trait, the operator `somewhatBigger` is defined to be equal to

```

biggerTrait: trait
  includes Integer
  introduces
    muchBigger, somewhatBigger: Int → Int
  asserts
    ∀ i: Int
      somewhatBigger(i) == muchBigger(i);
  implies
    converts somewhatBigger: Int → Int

```

Fig. 8. An LSL trait in which `somewhatBigger` is convertible, but `somewhatBigger(i)` is not completely-defined.

```

factTraitE: trait
  includes factTrait
  implies
    exact  $\forall k: \text{Int}$  such that  $k \geq 0$ 
      fact(k)

```

Fig. 9. A trait that demonstrates the **exact** clause. The **includes** directive has the effect of textually including the trait `factTrait` given above.

`muchBigger`; however, `muchBigger` is quite underspecified, since no assertions constrain it. Yet, the **converts** clause in the **implies** section is still provable, because `somewhatBigger` is completely-defined, relative to `muchBigger`. That is, once `muchBigger` is determined, `somewhatBigger` becomes completely-defined.

Because of this distinction between conversion and complete definition, we propose adding another implication clause to LSL. This clause, which we call the **exact** clause, has a form similar to that of the LSL **exempting** clause (although it would not be a subclass of a **converts** clause). The idea is that it would allow one to make redundant claims that terms are completely-defined. For example the **exact** clause in Fig. 9 says that terms of the form `fact(k)` are intended to be completely-defined, if $k \geq 0$. The syntax would be as follows.

```

exact-clause ::= 'exact' [quantifier] [such-that] term+,
such-that ::= 'such' 'that' term

```

The extra information in the **exact** clause, which does not affect the trait's theory, can be used to help debug an LSL specification, by trying to prove the following property.

Definition 3.1 (provable for exact clauses) Let T be a trait that contains an **exact** clause of the form **exact** $\forall \vec{a} : \vec{A}$ **such that** $Q(\vec{a}) P(\vec{a})$, where $Q(\vec{a})$ is a predicate and $P(\vec{a})$ is a term in the language of T . This clause is *provable for T* if and only if:

$$T \cup T' \vdash \forall \vec{a} : \vec{A}. (Q(\vec{a}) \wedge Q'(\vec{a})) \Rightarrow P(\vec{a}) = P'(\vec{a}) \quad (1)$$

For example, in Figure 9, the **exact** clause is provable for `factTraitE` if the following condition is provable from `factTraitE` \cup `factTraitE'`.

$$\forall k : \text{Int}. (k \geq 0 \wedge k \geq' 0) \Rightarrow \text{fact}(k) = \text{fact}'(k)$$

3.2. Exact Predicates

For use in proving protection, we define predicates of the form `Exact('E')`, based on the form (i.e., the text) of each expression E . These resemble the domain predicates, `Dom('E')`, described by some authors [Ger78, CoH79, Bli81]. However, they have a different purpose, since an operator, such as `choose` on nonempty sets, may be underspecified for a reason other than being partial. They also resemble the definedness predicate (D) used in studies of partial algebras [BrW82] and in COLD [FeJ92]; however D is defined model-theoretically, not syntactically. The definition of `Exact('')` is based on the **exact** clauses given in the trait's implications and those of included traits. This definition is lifted to arbitrary terms by requiring terms substituted for the variables in an **exact** clause to be themselves exact, and using the structure of terms formed from LSL's

$$\begin{aligned}
\text{Exact}('x') &= \text{true}, \text{ if } x \text{ is a variable} \\
\text{Exact}('P(\vec{E})') &= \bigwedge_{E_i \in \vec{E}} \text{Exact}('E_i') \wedge Q(\vec{E}), \\
&\quad \text{if the trait's } \mathbf{implies} \text{ section contains a clause:} \\
&\quad \mathbf{exact} \forall \vec{a} : \vec{A} \mathbf{such that} Q(\vec{a}) P(\vec{a}) \\
\text{Exact}('¬E') &= \text{Exact}('E') \\
\text{Exact}('E_1 \circ E_2') &= \text{Exact}('E_1') \wedge \text{Exact}('E_2'), \\
&\quad \text{if } \circ \text{ is } =, \neq, \text{ or a boolean operator: } \wedge, \vee, \text{ or } \Rightarrow \\
\text{Exact}('∀\vec{x} : \vec{T}. E') &= \forall \vec{x} : \vec{T}. \text{Exact}('E') \\
\text{Exact}('∃\vec{x} : \vec{T}. E') &= \forall \vec{x} : \vec{T}. \text{Exact}('E') \\
\text{Exact}('if } E_1 \mathbf{then } E_2 \mathbf{else } E_3') &= \text{Exact}('E_1') \\
&\quad \wedge \text{Exact}('E_2') \wedge \text{Exact}('E_3') \\
\text{Exact}('E') &= \text{false}, \text{ otherwise}
\end{aligned}$$

Fig. 10. Definition of Exact.

built-in trait operators (boolean operators, equality, and conditionals). See Fig. 10 for the definition.³

For example, for the trait of Fig. 9, the following holds.

$$\text{Exact}('fact(k)') = (k \geq 0)$$

3.3. Using Exact Predicates to Prove Underspec-Protection

Provided the information given in the **exact** clauses is provable for a trait T , then Exact predicates can be used as a sufficient condition for determining when a term is completely-defined for T .

Lemma 3.2. Let T be a trait in which each **exact** clause is provable for T . Let $R(\vec{x})$ be a term with free variables, $\vec{x} : \vec{U}$. If $T \vdash \forall \vec{x} : \vec{U}. \text{Exact}('R(\vec{x})')$, then $R(\vec{x})$ is completely-defined for T .

Proof. (By structural induction on terms.) Suppose $T \vdash \forall \vec{x} : \vec{U}. \text{Exact}('R(\vec{x})')$.

For the basis, suppose $R(\vec{x})$ is a variable x_i . Then $\forall \vec{x} : \vec{U}. x_i = x_i$ is trivially provable, and so x_i is completely-defined by definition.

For the inductive step, suppose that the result holds for all subterms of $R(\vec{x})$. If $R(\vec{x})$ is an invocation of some operator of T that is not a boolean operator, equality, inequality, or **if then else**, then by definition, it must be that $R(\vec{x})$ has the form $P(\vec{E}(\vec{x}))$ and that trait T has a clause of the form **exact** $\forall \vec{a} : \vec{A}$ **such that** $Q(\vec{a}) P(\vec{a})$. Furthermore, by definition of $\text{Exact}(' \cdot')$, it must be the case that

$$T \vdash \bigwedge_{E_i(\vec{x}) \in \vec{E}(\vec{x})} \text{Exact}('E_i(\vec{x})') \wedge Q(\vec{E}(\vec{x})) \quad (2)$$

Since T' is a primed copy of T , it must also be the case that

$$T' \vdash \bigwedge_{E'_i(\vec{x}) \in \vec{E}'(\vec{x})} \text{Exact}('E'_i(\vec{x})') \wedge Q'(\vec{E}'(\vec{x})) \quad (3)$$

Because the \vec{x} are free in the above two formulas, by universal generalization

³ The free variables of these terms are not important, so they are suppressed.

$$T \cup T' \vdash \forall \vec{x} : \vec{U}. Q(\vec{E}(\vec{x})) \wedge Q'(E'(\vec{x})) \quad (4)$$

By the inductive hypothesis, since each $E_i(\vec{x})$ is exact, for each i ,

$$T \cup T' \vdash \forall \vec{x} : \vec{U}. E_i(\vec{x}) = E'_i(\vec{x}) \quad (5)$$

Since the **exact** clauses are assumed to be provable for T , by definition we have

$$T \cup T' \vdash \forall \vec{a} : \vec{A}. (Q(\vec{a}) \wedge Q'(\vec{a})) \Rightarrow P(\vec{a}) = P'(\vec{a}) \quad (6)$$

Instantiating \vec{a} to $\vec{E}(\vec{x})$, we obtain the following.

$$T \cup T' \vdash \forall \vec{x} : \vec{U}. (Q(\vec{E}(\vec{x})) \wedge Q'(E'(\vec{x}))) \Rightarrow P(\vec{E}(\vec{x})) = P'(E'(\vec{x})) \quad (7)$$

Then using Formula (5), it follows that

$$T \cup T' \vdash \forall \vec{x} : \vec{U}. (Q(\vec{E}(\vec{x})) \wedge Q'(E'(\vec{x}))) \Rightarrow P(\vec{E}(\vec{x})) = P'(E'(\vec{x})) \quad (8)$$

But by (4), the hypothesis of this implication is provable, so it follows that $T \cup T' \vdash \forall \vec{x} : \vec{U}. P(\vec{E}(\vec{x})) = P'(E'(\vec{x}))$.

The other cases follow directly from the inductive hypothesis and the definition of $\text{Exact}(\cdot \cdot \cdot)$. \square

However, the converse to the above lemma does not hold. One reason is that the specifier of the used trait may not note when some terms are exact. But even if the information given is complete, the definition of Exact does not take into account other knowledge from the theory of the trait. For example, consider the trait `bufferTrait`, which is specified in Fig. 5. It specifies the constant `bufSize`, but `bufSize` is underspecified (hence no **exact** clause is given). The term

```
bufSize < 4096
```

is completely-defined for `bufferTrait`. However,

```
Exact('bufSize < 4096') = false,
```

because $\text{Exact}(\text{'bufSize'})$ is false.

Definition 3.3 (exact procedure specification) A procedure specification, S , that uses trait T , has formal parameters $\vec{x} : \vec{U}$, precondition $Q(\vec{x})$, and postcondition $R(\vec{x})$, is *exact* if and only if

- $T \vdash \forall \vec{x} : \vec{U}. \text{Exact}(\text{' } Q(\vec{x}) \text{'})$, and
- $T \vdash \forall \vec{x} : \vec{U}. Q(\vec{x}) \Rightarrow \text{Exact}(\text{' } R(\vec{x}) \text{'})$.

Our suggested technique for proving protection, therefore, is to prove that the specification in question is exact.

Corollary 3.4. Let T be a trait in which each **exact** clause is provable for T . Let S be a procedure specification that uses trait T . If S is exact, then S is underspec-protective.

Proof. Let $Q(\vec{x})$ be the precondition of S , and let $R(\vec{x})$ be its postcondition. Suppose S is exact. Then by definition, $T \vdash \forall \vec{x} : \vec{U}. \text{Exact}(\text{' } Q(\vec{x}) \text{'})$. So by Lemma 3.2, $Q(\vec{x})$ is completely-defined for T . By definition, $T \vdash \forall \vec{x} : \vec{U}. Q(\vec{x}) \Rightarrow \text{Exact}(\text{' } R(\vec{x}) \text{'})$. Suppose for each \vec{x} , $Q(\vec{x})$ holds; then for each \vec{x} , $\text{Exact}(\text{' } R(\vec{x}) \text{'})$ holds, and so by Lemma 3.2 $R(\vec{x})$ is completely-defined for T . \square

As an example of the use of the above corollary, we show how to prove that the specification of `factorial` in Fig. 7 is completely-defined with respect to the trait `factTraitE` of Fig. 9. To do this we prove that the specification is exact with respect to `factTraitE`. First, the precondition is exact, because $\text{Exact}('x \geq 0')$ is true. $\text{Exact}('0')$ is true, because 0 is a generator. We assume the trait `Integer` has been extended with implications that say that \geq is exact.

Then for the postcondition, one can calculate as follows, for all $x : \text{int}$.

```

x ≥ 0 ⇒ Exact('result = fact(x)')
= {by definition of Exact}
  x ≥ 0 ⇒ (Exact('result') ∧ Exact('fact(x)'))
= {by definition of Exact for fact}
  x ≥ 0 ⇒ (Exact('result') ∧ Exact('x') ∧ x ≥ 0)
= {by definition of Exact for variables, treating result as a variable}
  x ≥ 0 ⇒ (true ∧ true ∧ x ≥ 0)
= {by predicate calculus}
  true

```

However, if a procedure specification is protective, it is not necessarily exact. For example, a specification that uses the term `bufSize < 4096` as its precondition could be protective without being exact. Thus exactness is a sufficient, but not necessary, condition for protection.

4. Discussion of Underspec-Protection

One might wonder whether a procedure specification is underspec-protective if and only if it is deterministic. However, the two notions are independent. For example, the specification given in Figure 11 is protective (even exact) but very nondeterministic. It specifies a C++ procedure that can change the value of the object `x` (passed by reference) to any integer. Fig. 12 is an example of a specification that is not protective, because the precondition is not completely-defined, but the procedure specified must be deterministic when its precondition is met.

The notion of underspec-protection should also not be confused with the specification being “well-defined” in the sense of not containing any mathematically suspect terms. For example, the specification in Fig. 13 is not protective, because the operator `choose` is (intentionally) underspecified. However, the spec-

```

void chaos1(int& x);
behavior {
  modifies x;
  ensures true;
}

```

Fig. 11. The Larch/C++ specification of a procedure that is underspec-protective, even exact, but not deterministic. The formal parameter is passed by reference, because of the `&`.

```

uses bufferTrait;
int foo(int x);
behavior {
  requires bufSize < x;
  ensures result = 3;
}

```

Fig. 12. A specification that is deterministic but not underspec-protective.

```

uses IntSetTrait;
int pick(IntSet s);
behavior {
  requires size(s^) > 0;
  ensures result = choose(s^)  $\wedge$  s' = delete(choose(s^), s^);
}

```

Fig. 13. A specification that is “well-defined” but not underspec-protective. The notations s^{\wedge} and s' mean the starting and ending values of s .

Table 1. Summary of results related to underspec-protection.

Level	Facts	
Trait	exact \Rightarrow completely-defined completely-defined \neq convertible	Lemma 3.2 Fig. 8
BISL	exact \Rightarrow underspec-protective underspec-protective \neq deterministic well-defined $\not\Rightarrow$ underspec-protective	Corollary 3.4 Figs 11 and 12 Fig. 13

ification is well-defined because the precondition does protect `choose` from being applied outside its intended domain. (Note that the specification describes a set of functions that are each deterministic, but which individually can use any algorithm to pick elements of a set.) Thus a specification that is not protective is not necessarily ill-defined; there is no problem as long as the underspecification at the interface level is intentional.

Our technical results related to underspec-protection are summarized in Table 1. The main concept is determining when a procedure specification is protective, in the sense that it does not force implementations to satisfy unintended consequences of an LSL trait. We have given two proof techniques for proving protection. The first is equivalent to the definition and based on the notion of completely-defined terms. The second is a sufficient but not necessary test and based on the notion of exact terms, which makes it easier to apply. The concept of an exact term is based on an extension to LSL that allows one to specify which terms are not intended to be underspecified. This extension to LSL provides better documentation and allows enhanced debugging (in the sense of [GGH90] [GHG93, Chapter 7]) of LSL specifications.

5. Summary and Conclusions

In this paper we have given two definitions that are instances of the concept of protection. The definition of partiality-protection can be used with languages like VDM-SL and COLD-K, since these languages use a logic that admits the existence of partial functions. Underspec-protection is an analogous notion that is necessary for languages like Larch, RESOLVE, and Z, since they use logics that deal only with total functions.

Both kinds of protection may be useful in VDM-SL or COLD-K, where one can define partial functions and use underspecification. For example, after checking that a VDM-SL specification is partiality-protective, then one could check that it was also underspec-protective (assuming that the procedure was intended to be completely specified and not underspecified). Checks that a VDM-

SL procedure is underspec-protective can be done in same way as we described them for the Larch family.

Both kinds of protection may also be useful for writers of executable specifications. For example, in a language like Eiffel [Mey88], partiality-protection for a procedure would ensure that its precondition would be flagged as false instead of encountering an error, allowing an error to happen in its body, or encountering an error in its postcondition.

Acknowledgments

Thanks to Cliff Jones, Jim Horning, Adrian Fiech, Steve Garland, Clyde Ruby, Krishna Kishore Dhara, Matt Markland, and the anonymous referees for comments and discussions of earlier drafts. This paper is published with the understanding that the US Government is authorized to reproduce and distribute reprints for Governmental purposes, notwithstanding any copyright notation thereon.

Leavens's work was supported in part by NSF grant CCR-9593168, a faculty improvement leave from Iowa State, and the Computer Science and Engineering Department at the University of Washington.

Wing's research is sponsored by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency (ARPA) under grant number F33615-93-1-1330. Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of Wright Laboratory or the United States Government.

References

- [A⁺93] Andrews, D. J. et al.: Information technology programming languages — VDM-SL: First committee draft standard CD1387-1. Document ISO/IEC JTC1/SC22/WG19 N-20, International Standards Organization, November 1993. <ftp://gatekeeper.dec.com/pub/standards/vdmsl/>.
- [BCJ84] Barringer, H., Cheng, J. H., and Jones, C. B.: A logic covering undefinedness in program proofs. *Acta Informatica*, 21(3):251–269, October 1984.
- [BFL94] Bicarregui, J., Fitzgerald, J. S., Lindsay, P. A., Moore, R. and Ritchie, B.: *Proof in VDM: A Practitioner's Guide*. Springer-Verlag, New York, NY, 1994.
- [Bij90] Bijlsma, A.: Semantics of quasi-boolean expressions. In W. H. J. Feijen et al., editors, *Beauty is Our Business*, pages 27–35. Springer-Verlag, 1990.
- [Bli81] Blikle, A.: The clean termination of iterative programs. *Acta Informatica*, 16:199–217, 1981.
- [Bli91] Blikle, A.: Three-valued predicates for software specification and validation. *Fundamenta Informaticae*, XIV:387–410, 1991.
- [BrW82] Broy, M. and Wirsing, M.: Partial abstract types. *Acta Informatica*, 18(1):47–64, 1982.
- [CoH79] Coleman, D. and Hughes, J. W.: The clean termination of Pascal programs. *Acta Informatica*, 11:195–210, 1979.
- [ChJ90] Cheng, J. H. and Jones, C. B.: On the usability of logics which handle partial functions. In C. Morgan and J. C. P. Woodcock, editors, *Proceedings of the Third Refinement Workshop*, Workshops in Computing Series, pages 51–69, Berlin, 1990. Springer-Verlag.
- [FeJ92] Feijs, L. M. G. and Jonkers, H. B. M.: *Formal Specification and Design*, volume 35 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, UK, 1992.
- [GCD95] Gogolla, M., Conrad, S., Denker, G., Herzig, R., Vlachantonis, N. and Ehrig, H.: TROLL *light* — the language and its development environment. In Manfred Broy and Stefan Jähnichen, editors, *KORSO: Methods, Languages and Tools for the Construction*

- of Correct Software*, volume 1009 of *Lecture Notes in Computer Science*, pages 205–220. Springer-Verlag, New York, NY, 1995.
- [Ger78] German, S. M.: Automating proofs of the absence of common runtime errors. In *Conference record of the Fifth Annual ACM Symposium on Principles of Programming Languages*, pages 105–118. ACM, January 1978.
- [GGH90] Garland, S. J., Gutttag, J. V. and Horning, J. J.: Debugging Larch Shared Language specifications. *IEEE Transactions on Software Engineering*, 16(6):1044–1057, September 1990.
- [GHG93] Gutttag, J. V., Horning, J. J., Garland, S. J., Jones, K. D., Modet, A. and Wing, J. M.: *Larch: Languages and Tools for Formal Specification*. Springer-Verlag, New York, NY, 1993.
- [GHM90] Gutttag, J. V. Horning, J. J. and Modet, A.: Report on the Larch Shared Language: Version 2.3. Technical Report 58, Digital Equipment Corporation, Systems Research Center, 130 Lytton Avenue, Palo Alto, CA 94301, April 1990. Order from src-report@src.dec.com.
- [GrS95] Gries D. and Schneider, F. B.: Avoiding the undefined by underspecification. In Jan van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, number 1000 in *Lecture Notes in Computer Science*, pages 366–373. Springer-Verlag, New York, NY, 1995.
- [Hay93] Hayes, I. editor.: *Specification Case Studies*. International Series in Computer Science. Prentice-Hall, Inc., second edition, 1993.
- [JoM94] Jones, C. B. and Middelburg, K.: A typed logic of partial functions reconstructed classically. *Acta Informatica*, 31(5):399–430, 1994.
- [Jon90] Jones, C. B.: *Systematic Software Development Using VDM*. International Series in Computer Science. Prentice Hall, Englewood Cliffs, NJ, second edition, 1990.
- [Jon95] Jones, C. B.: Partial functions and logics: A warning. *Information Processing Letters*, 54(2):65–67, 1995.
- [KTB91] Konikowska, B., Tarlecki, A. and Blikle, A.: A three-valued logic for software specification and validation. *Fundamenta Informaticae*, XIV:411–453, 1991.
- [Lea97] Leavens, G. T.: Larch/C++ Reference Manual. Version 5.14. Available in ftp://ftp.cs.iastate.edu/pub/larchc++/lcpp.ps.gz or on the World Wide Web at the URL <http://www.cs.iastate.edu/~leavens/larchc++.html>, October 1997.
- [Mey88] Meyer, B.: *Object-oriented Software Construction*. Prentice Hall, New York, NY, 1988.
- [ORS95] Owre, S., Rushby, J., Shankar, N. and von Henke, F.: Formal verification for fault-tolerant architectures: Prolegomena to the design of PVS. *IEEE Transactions on Software Engineering*, 21(2):107–125, February 1995.
- [OSW94] Ogden, W. F., Sitaraman, M., Weide, B. W. and Zweben, S. H.: Part I: The RESOLVE framework and discipline — a research synopsis. *ACM SIGSOFT Software Engineering Notes*, 19(4):23–28, Oct 1994.
- [SKT95] Sannella, D., Kahrs, S. and Tarlecki, A.: The definition of Extended ML: a gentle introduction. Technical Report ECS-LFCS-95-322, Laboratory for Foundations of Computer Science, University of Edinburgh, October 1995. To appear in *Theoretical Computer Science*.
- [Spi92] Spivey, J. M.: *The Z Notation: A Reference Manual*. International Series in Computer Science. Prentice-Hall, New York, NY, second edition, 1992.
- [WDC95] Wolter, U., Didrich K., Cornelius F., Klar, M., R. Wessály, and H. Ehrig: How to cope with the spectrum of SPECTRUM. In Manfred Broy and Stefan Jähnichen, editors, *KORSO: Methods, Languages and Tools for the Construction of Correct Software*, volume 1009 of *Lecture Notes in Computer Science*, pages 173–189. Springer-Verlag, New York, NY, 1995.
- [Win83] Wing, J. M.: A two-tiered approach to specifying programs. Technical Report TR-299, Massachusetts Institute of Technology, Laboratory for Computer Science, 1983.
- [WoJ96] Woodcock, J. and Jackson, D. About the semantics of partial functions in Z. Personal communication, April 1996.

A. Appendix: Understanding Underspecification in LSL

A *partial function* is a function that does not give a value for some elements of its declared domain. For example, the operator that returns the head of a list can be modeled as a partial function on lists; if that is done, then `head(empty)` fails to denote an element. (That is, `head(empty)` is “undefined.”)

```

JonesExample1: trait
  includes Integer
  introduces
    it: → OneElem
    f: Int → OneElem
  asserts
    OneElem generated by it
    ∀ i: Int
      f(i) == if i=0 then it else f(i-1)
  implies
    converts f: Int → OneElem

```

Fig. 14. Jones’s first example, a function into a one-element set.

The logic used by the Larch Shared Language (LSL) [GHG93, Chapter 4] [GHM90] deals with partiality by using underspecification. As noted in the main body, this means that one avoids specifying a value for undefined terms, but the logic assumes that all functions are total. For example, `head(empty)` denotes some element of the appropriate type, even if the user has not specified what element that term denotes. Where an LSL specification is silent, terms take on some (unspecified) value.

In common with other logics that use underspecification to avoid the undefined [GrS95], the logic of LSL is classical, and thus has several pleasing formal properties. However, as Jones pointed out in a recent paper [Jon95], there are a few subtle aspects to this kind of logic that users should be aware of.

We translate Jones’s first example into the LSL trait shown in Fig. 14. This trait defines a sort, `OneElem`, a constant `it`, and a function `f`. Because of the **generated by** clause, the sort `OneElem` has only one element, the constant `it`. (The current version of LSL allows such sorts, contrary to [Jon95].) In LSL `f(-1) = it`, because `f` has to take on some value when applied to `-1`, and the only possible value is `it`. Although Jones notes that this is “not an inconsistency” he says that “it is certainly likely to surprise someone who views” the definition of `f` as specifying “a partial function” (p. 66). Another way of putting Jones’s point is that it is simply impossible to specify partial functions in LSL, even using recursion.

Jones’s other major example brings out a more important warning about the underspecification approach. This example is a recursive definition of the factorial function, and is translated into LSL in Fig. 15. Jones’s warning about this example is that, in a logic such as LSL’s, a model of fact must satisfy irrelevant equations such as the following, which is also highlighted in the redundant **implies** section of the trait.

$$\text{fact}(-1) == -\text{fact}(-2) \tag{9}$$

This follows because `fact(-1)` denotes some (unspecified) value.

Jones’s warning could have been stated more strongly, since not only is there a danger that one might specify unwanted properties, but there is also a danger that these unwanted properties might cause inconsistency. The trait `factTrait` of Fig. 15, actually has quite a few such unwanted equations but manages to escape inconsistency because of special properties of the integers. (That is, the following equations are also consequences of the trait.

$$\begin{aligned} \text{fact}(-1) &== (-1) * \text{fact}(-2) \\ \text{fact}(-2) &== (-2) * \text{fact}(-3) \\ \text{fact}(-3) &== (-3) * \text{fact}(-4) \end{aligned}$$

```

factTrait: trait
  includes Integer
  introduces
    fact: Int → Int
  asserts
    ∀ i: Int
      fact(i) == if i=0 then 1 else i * fact(i-1);
  implies
    equations
      fact(3) == 6;
      fact(-1) == - fact(-2);

```

Fig. 15. Jones's factorial example.

```

badRecTrait: trait
  includes Integer
  introduces
    zero: Int → Int
  asserts
    ∀ k: Int
      zero(k) == if k = 0 then 0 else min(k, zero(k-1));
  implies
    ∀ k : Int
      zero(-1) < -1;
      zero(-1) < k;

```

Fig. 16. A trait with an inconsistent recursive definition.

However, the trait is not inconsistent, because one can let $\text{fact}(i) == 0$ for all negative integers i , which allows all these equations to be satisfied.)

To illustrate what can happen if one is not careful, consider the trait in Fig. 16, `badRecTrait`. At first glance, it looks like `zero` is a (silly) definition of a constant function that returns zero for any nonnegative integer. However, this specification is not careful to explicitly underspecify the value of `zero` for negative arguments. That is, although the specifier might think that it does not matter what `zero` returns for negative arguments, just ignoring the issue in the specification does not mean that the value is underspecified. For example, what does the specification say about `zero(-1)`? It is easy to see that it is less than -1 , and less than -2 , and indeed less than any integer. But the `Integer` trait in Guttag and Horning's handbook (see [Guttag-Horning93], p. 163) does not allow there to be such an integer; so this trait is inconsistent, because the value of `zero(-1)` is *overspecified* — it has to satisfy too many constraints. Although none of these constraints were intended, the trait is just as inconsistent as if they were intentionally specified.

To avoid the possibility of such inconsistency arising from unintentional overspecification, it is best to use intentional underspecification. That is, to avoid the possibility that an operator may be inconsistently specified (and the need to prove that the inconsistency does not happen), it is best to use conditional equations instead of unguarded recursive equations. For example, one can write `factTrait` as in Fig. 6, where the equation for the recursive case is only postulated to hold for its intended domain [GrS95]. By writing `factTrait` in that way, one avoids postulating Equation (9); that is, nothing at all is specified about the value of `fact(-1)`.

B. Appendix: Conversion in LSL

This appendix explains the notion of conversion in LSL.

In an LSL trait, one can state redundant properties (theorems) that one believes do (or should) hold. These redundant properties are stated in the **implies** section of the specification. Proofs of such properties can be attempted, and are a way of debugging the trait [GGH90] [GHG93, Chapter 7].

For our purposes, the most interesting kind of redundant property one can state in the **implies** section is that an operator is well-defined with respect to other operators. This is done by using a **converts** clause, as was done in Fig. 14. A **converts** clause says that the axioms of the trait uniquely define the operators named in the clause, “relative to the other operators in the trait” [GHG93, p. 142]. To prove this, one must show it for all possible arguments. The Larch Prover (LP) uses the following proof technique [GHG93, pp. 142–4]. Let $T(\vec{f})$ be a trait, which names operators \vec{f} in **converts** clauses in its **implies** section. Let $T(\vec{f}')$ be a version of the trait $T(\vec{f})$ in which each of the operators f_i named in a **converts** clause is replaced by f'_i . Then one proves, for each such $f_i : \vec{A} \rightarrow B$,

$$T(\vec{f}) \cup T(\vec{f}') \vdash \forall \vec{a} : \vec{A}. f_i(\vec{a}) = f'_i(\vec{a}) \quad (10)$$

The proof would show that there cannot be two different interpretations of the operator f_i .

For example, to prove the **converts** clause for \mathbf{f} in Fig. 14, one axiomatizes an operator \mathbf{f}' in the same way as \mathbf{f} , and then proves the following.

$$\forall i: \text{Int} \quad \mathbf{f}(i) == \mathbf{f}'(i)$$

(This is proved by using the rule given by the **generated by** clause in Fig. 14.)

Often one wants to prove that an operator is converted, except for some arguments. For example, one would want to prove that the head operator on lists is converted, except that $\text{head}(\text{empty})$, which is purposely left underspecified. To do this one uses a **converts** clause of the following form in LSL.

```
converts
  head: List[T] → T
  exempting head(empty)
```

The **exempting** clause allows the specifier to state what terms are intentionally underspecified. In terms of the proof that head is converted, except where it is not intentionally underspecified, the exempting clause allows one to use the following equation

$$\text{head}(\text{empty}) == \text{head}'(\text{empty})$$

in the proof that, for all lists l , $\text{head}(l) == \text{head}'(l)$.

Received October 1997

Accepted in revised form March 1998 by C. B. Jones