

Hints to Specifiers

Jeannette M. Wing

May 22, 1995

CMU-CS-95-118R

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This paper is derived from the paper “Teaching Mathematics to Software Engineers,” which will appear in the *Proceedings of the Fourth International Conference on Algebraic Methodology and Software Technology '95*. The AMAST paper is the basis of an invited talk for AMAST’s Education Day, July 3, 1995, in Montreal, Canada.

Abstract

I present a list of hints for writing specifications. I address high-level issues like learning to abstract and low-level issues like getting the details of logical expressions right. This paper should be of interest not only to students of formal methods but also to their teachers.

This 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.

Keywords: formal methods, formal specification, software engineering, education, discrete mathematics, mathematical logic, algebraic specifications, Z, Larch

1. Motivation

Over the years I have been accumulating hints that I give students in response to common problems and recurrent questions that arise as they try their hand at writing specifications. I often remind myself of these hints when I write specifications too. I've broadly categorized them along the following dimensions:

- Figuring out *why* you are going through this specification effort (Section 2). What do you hope to get out of using formalism?
- Figuring out *what* of the system you want *to specify* (Section 3).
- Figuring out *how* to specify (Section 4). The most important hurdle to overcome is learning to abstract. I also give specific suggestions on how to make incremental progress when writing a specification.
- Figuring out *what to write down* (Section 5). Learn and abide by a formal method's set of conventions but do not feel unduly constrained by them. Also, we all make logical errors sometimes; I point out some common troublespots in getting the details of a specification right.

My hints are targeted for the novice specifier, but experts, such as teachers of formal methods, may also find them useful.

I will illustrate my points with examples, usually in Z or Larch. Many actually make more than one point.

2. Why Specify?

You should first ask yourself this question, "Why specify?" You might choose to specify because you want additional documentation of your system's interfaces, you want a more abstract description of your system design, or you want to perform some formal analysis of your system. What you write should be determined by what it is you want to do with your specification.

You should then ask yourself "Why *formally* specify?" Your answer determines what is to be formalized, what formal method to use, and what benefits you expect from a formal specification not attainable from an informal one. When I have asked this question of system builders, here are the kinds of responses I have heard:

- Showing that a property holds globally of the entire system.
 - I want to characterize the "correctness condition" I can promise the user of my system.
 - I want to show this property is really a system invariant.
 - I want to show my system meets some high-level design criteria.
- Error handling
 - I want to specify what happens if an error occurs.

- I want to specify the right thing happens if an error occurs.
 - I want to make sure this error never occurs.
- Completeness
 - I want to make sure that I’ve covered all the cases, including error cases, for this protocol.
 - I’d like to know that this language I’ve designed is computationally complete.
- Specifying interfaces.
 - I’d like to define a hierarchy of C++ classes.
 - I’d like a more formal description of this system’s user interface.
- Getting a handle on complexity.
 - The design is getting too complicated. I can’t fit it all in my head. I need a way to think about it in smaller pieces.
- Change control.
 - Every time I change one piece of code I need to know what other pieces are affected. I’d like to know where else to look without looking at all modules, without looking at all the source code.

Judicious use of formalism can help address all these problems to varying levels of detail and rigor.

3. What to Specify?

Formal methods are not to the point where an entire large, software system can be completely specified. You may be able to specify one aspect of it, e.g., its functionality or its real-time behavior; you may be able to specify many aspects of a part of it, e.g., specifying both functionality and real-time behavior of its safety-critical part. In practice, you may only care to specify one aspect of a part of a system anyway.

In writing a specification, you should decide whether it is describing *required* or *permitted* behavior. Must or may? Since a specification can be viewed as an abstraction of many possible, legitimate implementations, you might most naturally associate a specification with describing permitted behavior. An implementation may have any of the behaviors permitted by the specification, but the implementor is not required to realize all. For example, a nondeterministic *choose* operation specified for sets will have a deterministic implementation. However, the expression “software system requirements” suggests that a customer may in fact *require* certain behavior. For example, in specifying an abstract data type’s interface, the assumption is that all, not some proper subset, of the operations listed must be implemented.

Once it is clear what you want from the specification process, you can turn to determining exactly what should be formalized.

In increasing order of level of detail, you might want to formalize a global *correctness condition* for the system, one or more system *invariants*, the *observable behavior* of a system, or *properties* of entities in a system.

Correctness Conditions

You usually have some informal notion of a global *correctness condition* that you expect your system to maintain. It might be something as standard as serializability, cache coherence, or deadlock freedom. Or, it might be very specific to the protocol or system at hand. If it is standard, then very likely someone else has developed a formal model for characterizing a system and a logic within which the correctness condition can be formally stated and proved. E.g., serializability has been thoroughly studied by the database community from all angles, theoretical to practical. If your correctness condition can be cast in terms of a well-known theory, it pays to reuse that work and not invent from scratch.

If it is not standard then an informal statement of the correctness condition should drive the formalization of the system model and expression of the correctness condition. For example, in work by Mummert et al. [14], the authors started with this informal statement of cache coherence for a distributed file system:

If a client believes that a cached file is valid then the server that is the authority on that file had better believe that the file is valid.

They developed a system model (a state machine model) and logic (based on the logic of authentication [5]) that enabled them to turn the informal statement into the following formal statement:

For all clients C , servers S , and objects d for which S is the repository,
if C **believes** *valid*(d_C) then S **believes** *valid*(d_C).

where clients, servers, objects, repository, **believes** and *valid* are formally defined concepts. The point is that the formal statement does not read too differently from the informal one.

Keep in mind this rule-of-thumb when formalizing from an informal statement: Let the things you want to describe formally drive the description of the formal model. There is a tendency to let the formal method drive the description of the formal model; you end up specifying what you can easily specify using that method. That is fine as far as it goes, but if there are things you cannot say or that are awkward to express using that method, you should not feel bound by the method. Invent your own syntax (to be defined later), add auxiliary definitions, or search for a complementary method.

The process of constructing a formal model of a system and formally characterizing the intended correctness condition can lead to surprises. More than once I have seen Ph.D. students start formalizing the systems that they were building and then have to back off from their expected and desired correctness condition. They end up realizing that it was too strong, not always guaranteed (e.g., not guaranteed for some failure case or for a “fast-path” case), or only locally true (holds for a system component but not the entire system). Correctness conditions for distributed systems are likely to be weaker than expected or desired because of the presence of failures (nodes or

links crashing) and transmission delays; the time to recover from failures and the time to transmit messages introduce “windows of vulnerability” during which the correctness condition cannot be guaranteed.

Invariants

The most common way to characterize certain kinds of correctness conditions is as a *state invariant*. An invariant is a property that does not change as the system goes from state to state. Remember also:

- An invariant is just a predicate. Given an appropriate assertion language, it is usually not a big deal to express an invariant formally.
- “True” is an invariant of any system. It’s the weakest invariant and hence not a very useful one; you probably want to say something more interesting about your system. If “true” ends up being your strongest invariant, revisit your system design.
- An invariant can serve multiple purposes. It is usually used to pare down a state space to the states of interest. For example, it can be used to characterize the set of *reachable* states or the set of *acceptable/legal* (“good”) states. (These two sets are not always the same. For example, you might want the set of acceptable states to be a subset of the set of reachable ones.) *Representation invariants* are used to define the domain of an abstraction function, used when showing that one system “implements” another [11].
- Different formal methods treat invariants differently. (See *Implicit versus Explicit* in Section 5.1 for an elaboration of this point.) Make sure you understand invariants in the context of the formal method you are using.
- Hard questioning of system invariants can lead to radically new designs.

To illustrate the last point, consider this example from the garbage collection community. One class of copying garbage collection algorithms relies on dividing the heap into two semi-spaces, *to-space* and *from-space*; in one phase of these algorithms, objects are copied from from-space to to-space [2]. Traditional copying garbage collection algorithms obey a “to-space invariant”: The user accesses objects only in to-space. Nettles and O’Toole observed that breaking this invariant and maintaining an alternative “from-space invariant” (the user accesses objects only in from-space) leads to simpler designs that are much easier to implement, analyze, and measure [15]. This observation led to a brand new class of garbage collection algorithms.

Observable Behavior

State invariants are a good way to characterize desired system properties. Formalizing state transitions will allow you to prove that they are maintained. When you specify state transitions, what you are specifying is the behavior of the system as it interacts with its environment, i.e., the system’s observable behavior.

It might seem obvious that what you want to specify is the observable behavior of a system, but sometimes when you are buried in the details of the task of specifying, you forget the bigger picture. Suppose you take a state machine approach to modeling your system. Here is a general approach to specifying observable behavior:

1. Identify the level of abstraction (see Section 4.1) at which you are specifying the system. This level determines the interface boundary that you are specifying; it determines what is or is not observable. For example, a bus error at the hardware level is not expected to be an observable event in the execution of an text formatter like `Word`, but `core dumped` is certainly an observable event when using a text editor like `emacs`.
2. Characterize the observable entities in a state at that given level of abstraction. These entities are sometimes called a system's *state variables* or *objects*. This step forces you to identify the relevant abstract *types* of your system (See the section on *Properties of State Entities* below.)
3. Characterize a set of initial states, and if appropriate, a set of final states.
4. Identify the operations that can access or modify the observable entities. These define your state transitions.
5. For each operation, characterize its observable effect on the observable state entities. For example, use Z schemas, Larch interfaces, or VDM pre/post-conditions.

Observable behavior should include any change in state that is observable to the user. If you are specifying an operation, then the kinds of observable state changes include changes in value to state entities, observable changes in the store (new entities that appear and old entities that disappear), results returned by the operation, and signaled exceptions or errors.

Another way to think about observable behavior is to think about *observable equivalence* [13, 10]. Ask “Can I distinguish between these two things?” where “things” might be states, individual entities in a state, traces of a process, or behavior sets of a process, depending on what you are specifying. If the answer is “yes,” then there must be way to tell them apart (perhaps by using unique names or perhaps by defining an *equal* operation); if “no,” then there must not be any way for the observer to tell them apart.

Properties of State Entities

The most important property to express of any entity in a system is its *type*. This statement is true regardless of the fine distinctions between the different type systems that different formal methods and specification (and programming) languages have. Since for a specification we are not concerned about compile-time or run-time costs of checking types, there is never a cost incurred in documenting in a specification an entity's type.

Since a type can be viewed as an abbreviation for a little theory, declaring an entity's type is a succinct way of associating a possibly infinite set of properties with the entity in one or two words. A truly powerful abstraction device!

For entities that are “structured” objects (e.g., an object that is a collection of other objects), when determining its type, the kinds of distinguishing properties include:

- Ordering. Are elements ordered or unordered? If ordering matters, is the order partial or total? Are elements removed FIFO, LIFO, or by priority?
- Duplicates. Are duplicates allowed?

- Boundedness. Is the object bounded in size or unbounded? Can the bound change or it is fixed at creation time?
- Associative access. Are elements retrieved by an index or key? Is the type of the index built-in (e.g., as for sequences and arrays) or user-definable (e.g., as for symbol tables and hash tables)?
- Shape. Is the structure of the object linear, hierarchical, acyclic, n -dimensional, or arbitrarily complex (e.g., graphs, forests)?

For entities that are relations, the kinds of distinguishing properties include whether the relation is a function (many-to-one), partial, finite, defined for only a finite domain, surjective, injective, bijective, and any (meaningful) combination of these.

Finally, algebraic properties help characterize any relational entity or any function or relation defined on a structured entity. The standard algebraic properties include: idempotency, reflexivity, symmetry, transitivity, commutativity, associativity, distributivity, existence of an identity element, and existence of an inverse relation or function. Algebras are well-known mathematical models for abstract data types and for processes [8, 13, 1]. For example, this algebraic equation characterizes the idempotency of inserting the same element into a set multiple times:

$$\text{insert}(\text{insert}(s, e), e) = \text{insert}(s, e)$$

and this characterizes *insert*'s commutativity property:

$$\text{insert}(\text{insert}(s, e_1), e_2) = \text{insert}(\text{insert}(s, e_2), e_1)$$

It also makes sense to ask about whether algebraic properties hold for operations on processes. For example, for CSP processes, parallel composition is both commutative and associative:

$$\begin{aligned} P \parallel Q &= Q \parallel P \\ P \parallel (Q \parallel R) &= (P \parallel Q) \parallel R \end{aligned}$$

4. How to Specify?

Given that you understand why you are specifying and what it is you want to specify, in what ways should you try to think about the system so that you can begin to specify and then make progress in writing your specification? The fundamental techniques are *abstraction* and *decomposition*. In specifying large, complex systems, abstraction is useful for focusing your attention to one level of detail at a time; decomposition, for one small piece of the system (at a given abstraction level) at a time. Both enable local reasoning.

4.1. Learn to Abstract: Try Not to Think Like a Programmer

The skill that people find the most difficult to acquire is the ability to abstract. One aspect of learning to abstract is being able to think at a level higher than programmers are used to.

Try to think definitionally not operationally.

A student said the following to me when trying to explain his system design:

If you do this and then that and then this and then that, you end up in a good state.

But if you do this and then that and then this, you end up in a bad state.

When specifying concurrent systems, rather than thinking of what characterizes all good states, people often think about whether a particular sequence of operations leads to a good or bad state. Taking this operational approach means ending up trying to enumerate all possible interleavings; this enumeration process quickly gets out of control, which is typically when a student will come knocking at my door for help. This problem is related to understanding invariants (see *Invariants* in Section 3). Invariably, the very first thing I need to teach students when I work with them one-on-one is what an invariant is.

*Try not to think computationally.*¹

When writing specifications, abstraction is intellectually liberating because you are not bound to think in terms of computers and their computations.

The following predicate

$$s = s' \wedge \langle e \rangle$$

might appear in the post-condition of the specification of a *remove* operation on sequences. Here, s stands for the sequence's initial value; s' , its final value; e , the element removed and returned. You most naturally might read “=” as assignment (especially if you are a C programmer) and not as a predicate symbol used here to relate values of objects in two different states. You may need to stare at such predicates for a while before realizing the assertional nature (and power) of logic.

Try constructing theories, not just models.

Building models is an abstraction process; but defining a theory takes a different kind of abstraction skill. When you construct a model of a system in terms of mathematical structures like sets, sequences, and relations, you get all properties of sets and relations “for free.” This has the advantage that you do not have to spell them out every time you specify a system, but the disadvantage that some of those properties are irrelevant to your system. Thus, in a model-based constructive approach, you also need to provide a way to say which properties about the standard mathematical structures may be irrelevant. (You might strip away of properties by using invariants.) For example, you might specify a stack in terms of a sequence, where the top of the stack corresponds to one end of the sequence. Then, you need not only to state which end of the sequence serves as the top of the stack, but also to eliminate some sequence properties, e.g., being able to index into a sequence or concatenate two sequences, because they have no relevance for stacks.

By contrast, in a theory-based approach you state explicitly exactly what properties you want your system to have. Any model that satisfies that theory is deemed to be acceptable. For example, the essence of stacks is captured by the well-known equations:

¹Another way of saying the same thing as above.

$$\begin{aligned} \text{pop}(\text{push}(s, e)) &= s \\ \text{top}(\text{push}(s, e)) &= e \end{aligned}$$

Sequences, or any other data structure, do not enter the picture at all.

Like many, you may find methods like Z and VDM appealing because they encourage a model-based rather than theory-based approach to specification. You can build up good intuition about your system if you have a model in hand. However, to practice learning how to abstract, try writing algebraic or axiomatic assertions about the model.

4.2. How To Proceed: Incrementally

At any given level of abstraction, we ignore some detail about the system below. You might feel anxious to specify everything for fear of being “incomplete.” Learning to abstract means learning when it is okay to leave something unspecified. This aspect of the abstraction process also allows incremental specification. In general, it is better to specify something partially than not at all.

Here are four common and important examples of incremental abstraction techniques: (1) first assume something is true of the input argument and capture this assumption in a pre-condition, then weaken the pre-condition; (2) first handle the normal case, then the failure case; (3) first ignore the fact that ordering (or no duplicates, etc.) matters, then strengthen the post-condition; (4) first assume the operation is atomic, then break it into smaller atomic steps. Let’s look in turn at each of these examples in their generality and in more detail.

Use pre-conditions.

Putting your “programmer’s” cap on, think of pre-conditions in the context of procedure call. A pre-condition serves two purposes: an obligation on the caller to establish before calling the procedure and an assumption the implementor can make when coding the procedure.

More generally, pre-conditions are a way of specifying assumptions about the environment of a system component. Such assumptions can and should be spelled out and written down explicitly. By doing so, you can specify and reason about a piece of the system without having to think about the entire system all at once. Thus pre-conditions assist in partial specification, incremental design, and local reasoning—all attractive means of dealing with the complexity of large software systems.

One technical difficulty that trips some people is what a specification means if a pre-condition is not met. In many specification techniques (like Z and Larch), when an operation’s pre-condition is not met, the interpretation is “all bets are off.” The interpretation is that the pre-condition is a *disclaimer*.² In other words, the operation is free to do anything, including not terminate, if the pre-condition does not hold. The technical justification is that when an operation is specified using pre- and post-conditions, the logical interpretation of the specification is an implication:

$$pre \Rightarrow post$$

²Thanks to Daniel Jackson for this term.

When the pre-condition is “false” then the implication is vacuously true, so any behavior should be allowed.

Some formal methods (like InaJo [17] and I/O automata [12]) use the term “pre-condition” but mean something entirely different. The pre-condition is interpreted as a *guard*; no state transition should occur if the guard is not met. Here the interpretation is conjunction:

$$pre \wedge post$$

The difference is that under the disclaimer interpretation, for any state s in which the pre-condition does not hold, the state pair, $\langle s, s' \rangle$, for any state s' , would be in the state transition relation; under the guard interpretation, no such state pair would be in the relation.³

There are other possible interpretations: For example, if the pre-condition is not met, it could mean that the state transition always goes to a special “error” state and termination is guaranteed, or it could mean the state transition leads to either an “error” state or non-termination. The point is that you must understand in the notation you are using what it means when a pre-condition is met or not met.

Finally, in the presence of concurrency, you need to specify both kinds of conditions for an operation: a pre-condition (as a disclaimer) and a guard. The pre-condition is evaluated in the state in which the operation is called; the guard, in the state in which the operation begins executing. Because of concurrency, a scheduler may delay the start of the execution of an operation to some time after the call of the operation; since there is time between the state in which the operation is called and the state in which it starts executing, an intervening operation (executed by some other process) may change the system’s state. Thus, a predicate that holds in the state when the operation is called may no longer hold in the state when the operation begins to execute. The point is to realize that in the presence of concurrency, there is a new kind of condition to specify.⁴

Specify errors/exceptions/failures.

It is as important to specify erroneous or exceptional behavior as it is to specify normal behavior. If an operation can lead to an undesired state, you should specify the conditions under which this state is reachable. If you are lucky, the specification language has some notational convenience (e.g., Larch’s **signals** clause) or prescribed technique (e.g., Z’s schema calculus) to remind you to describe error conditions; otherwise, handling errors may have to be disguised in terms of input or output arguments that serve as error flags.

There is a close correlation between pre-conditions and handling errors. Z specifiers draw this connection by abiding by this convention using schema disjunction:

$$TotalOp = NormalOp \vee ErrorOp$$

³There is further confusion in understanding pre-conditions in Z because even though you might write explicitly in your schema the conjunction, $xpre \wedge post$, where $xpre$ is the “explicit” pre-condition, the meaning is the implication, $pre \Rightarrow post$, where pre is the calculated pre-condition and usually not identical to $xpre$ [6].

⁴Larch calls the guard a *when-condition* to distinguish it from the standard pre-condition written in a **requires** clause.

where `NormalOp` is the specification (schema) of the `Op` operation under normal conditions, and `ErrorOp` is the specification of `Op` under the condition in which the pre-condition (which must be *calculated* [18] from `NormalOp`) does not hold. Thus, `TotalOp` gives the specification of `Op` under all possible conditions.

Larch specifiers, on the other hand, draw the connection by weakening the pre-condition, e.g., defining it to be equivalent to “true,” and correspondingly strengthening the post-condition. Thus,

```
Op = op()
    requires P
    ensures Q
```

turns into:

```
Op = op() signals (error)
    requires true
    ensures if P then Q else signal error
```

For interfaces to distributed systems, you cannot ignore the possibility of failure due to network partitions or crashed nodes. You could abstract from the different kinds of failures by introducing a generic “failure” exception that stands for errors arising from the distributed nature of your system.

The two main points to remember are (1) in support of incremental specification, specify the normal case and then handle the error cases, but (2) do not forget to handle the error cases!

Use nondeterminism.

Introducing nondeterminism is an effective abstraction technique. Nondeterminism permits design freedom and avoids implementation bias.

Nondeterminism may show up in many ways. It may be inherent to the behavior of an operation or object. Consider the *choose* operation on sets:

```
choose = op (s: set) returns (e: elem)
    requires s ≠ ∅
    ensures e ∈ s
```

The post-condition says that the element returned is a member of the set argument; it does not specify exactly which element is returned.

You can express nondeterminism by explicit use of disjunction in a post-condition:

```
traffic_light = op() returns (c: color)
    ensures c = red ∨ c = amber ∨ c = green
```

If the type `color` ranges over *red*, *amber*, *green*, and *blue*, the use of negation allows you to express the same property more succinctly:

```
traffic_light = op() returns (c: color)
    ensures c ≠ blue
```

You can express nondeterminism by explicit use of an existential quantifier, which is the more

general case of disjunction:

```
positively_random = op () returns (i: int)
ensures  $\exists x : int . i > x$ 
```

From a state machine model viewpoint (for instance when discussing deterministic and nondeterministic finite state automata), nondeterminism should not be confused with choice. Suppose δ is a state transition relation,

$$\delta : State, Action \rightarrow 2^{State}$$

Then an example of choice is:

$$\begin{aligned}\delta(s, a_1) &= \{t\} \\ \delta(s, a_2) &= \{u\}\end{aligned}$$

which says from state s you can either do the action a_1 (and go to the next state t), or do the action a_2 (and go to the next state u). However, an example of nondeterminism is:

$$\delta(s, a_1) = \{t, u\}$$

which says from state s you can do action a_1 and go to either state t or u .

Some formal methods for concurrent systems introduce their own notions of nondeterminism/choice; for example, CSP has two operators, one for internal choice (\sqcap) made by the machine and the other for external choice (\sqbox) made by the environment⁵. CCS has yet a different way to model nondeterminism.

The two main points are that (1) nondeterminism is a useful and important way to abstract, but (2) be careful to understand any given method's way of modeling nondeterminism/choice to use it properly.

Use Atomic Operations

For any system it is important to identify what the *atomic* operations are. An atomic operation is one whose execution is indivisible; only the states before and after its execution are observable. At any level of abstraction an atomic operation may be implemented in terms of sequences of lower-level atomic operations (e.g., a *write* operation to a file on disk might be implemented in terms of a sequence of *write* operations to individual disk blocks). Even assignment can be broken down into sequences of loads and stores to/from memory and registers.

It is usually assumed that each procedure of a sequential program is executed atomically; this assumption is rarely stated explicitly.

For a concurrent system, it is critical to state explicitly what operations are atomic. The atomicity of an operation, Op , guarantees that no other operation can interfere with Op 's execution and that you can abstract away from any intermediate (lower-level) state that it might actually pass through.

⁵Hoare calls the former “nondeterministic or” and the latter “general choice.”

5. What to Write?

With your pen poised over a blank sheet of paper or fingers over your keyboard, you now face the problem of what to write down. If you are using a specific formal method like Z, VDM, or Larch, you must know the syntax and semantics of its specification language. It is not enough to know what the syntactic features are; you need to understand what each means.

It is important to understand the difference between syntax and semantics. For example, a typical algebraic specification language has grammatical rules for formulating syntactically legal terms out of function and variable symbols. Each syntactically legal term denotes a value in some underlying algebraic model. For example, the term $insert(insert(\emptyset, e_1), e_2)$ is a syntactic entity that denotes the set value $\{e_1, e_2\}$, which is a semantic entity. For a standard model of sets, the syntactically different term $insert(insert(\emptyset, e_2), e_1)$ denotes the same semantic set value.

Associated with any formal method is its *assertion language*, usually based on some variation of first-order predicate logic. With assertions you nail down precisely your system's behavior. It is in your assertions where the smallest change in syntax can have a dramatic change in semantics. Getting the details of your assertions right is typically when you discover most of the conceptual misunderstandings of your system's design.

5.1. General Rules-of-Thumb

What distinguishes a formal method from mathematics is its methodological aspects. A specification written in the style of a given formal method is usually not just an unstructured set of formulae. Syntactic features make it easier to read the specification (e.g., the lines in a Z schema), remind the specifier what to write (e.g., the **modifies** clause in Larch), and aid in structuring a large specification into smaller, more modular pieces (e.g., Z schemas, Larch traits).

Implicit vs. Explicit

Most formal methods have well-defined specification languages so the choice of what you explicitly write down is guided by the grammar and constructs of the language.

However, there is a danger of forgetting the power of the unsaid. What is *not* explicitly stated in a specification often has a meaning. A naive specifier is likely to be unaware of these implicit consequences, thereby be in danger of writing nonsense. Here are three examples.

The first example is the *frame* issue. If you are specifying the behavior of one piece of the system in one specification module, you should say what effects that piece has *on the rest* of the system. In some formal methods (e.g., InaJo), you are forced to say explicitly what other pieces of the system do not change (NC''):

$$NC''(x_1, \dots, x_n)$$

This is sometimes impractical if n is large, or worse, if you do not or cannot know what the x_1, \dots, x_n are in advance.

In some methods (e.g., Larch), you say only what may (but is not required to) change; anything not listed explicitly is required *not* to change:

modifies y_1, \dots, y_m

This says $y_1 \dots y_m$ may change but the rest of the system stays the same.

A subtler point about the Larch **modifies** clause is that there is significance to the omission of the clause. The absence of a **modifies** clause says that no objects may change. Thus, if you write a post-condition that asserts some change in value to an input argument or global, the assertion would be inconsistent with an omitted **modifies** clause.

Z's Δ and Ξ operators on schemas are similar to InaJo's *NC* construct; they allow you to make statements local to individual operations about whether they change certain state variables or not. Use of these schema operators on say the schemas, S_i , leaves implicit the invariant properties of the system captured in S_i . These properties can be made explicit by “unrolling” the schemas S_i .

This feature of Z is related to my second example of implicit vs. explicit specification, which has to do with invariants. In some formal methods like Z, state invariants are stated explicitly. They are a critical part of the specification, i.e., the “property” component of a Z schema, and used to help calculate operation pre-conditions. In others like Larch, they are implicit and must be proved, usually using some kind of inductive proof rule. Finally, in others like the 1980 version of VDM [9], they are redundant. They are stated explicitly and contribute to the checklist of proof obligations generated for each operation.

Finally, the third example has to do with implicit quantification. In many algebraic specification languages the i equations in this list

$$\begin{array}{l} e_1 \\ \vdots \\ e_i \end{array}$$

are implicitly conjoined and quantified as follows:

$$\exists f_1 \dots \exists f_n . \forall x_1 \dots \forall x_m . e_1 \wedge \dots \wedge e_i$$

where $f_1 \dots f_n$ are the function symbols and $x_1 \dots x_m$ are the variables that appear in $e_1 \dots e_i$.

This kind of implicit quantification has subtle consequences. Consider the following (incorrect) equational specification of an operation that determines whether one set is a subset of another:

$$s_1 \subseteq s_2 = (e \in s_1 \Rightarrow e \in s_2)$$

What you really mean is:

$$s_1 \subseteq s_2 = \forall e . (e \in s_1 \Rightarrow e \in s_2)$$

but in most algebraic specification languages, writing a quantifier in the equation is syntactically illegal; the tipoff to the error is the occurrence of the free variable e on the righthand side of the first equation.

Auxiliary Definitions

Do not be afraid to use auxiliary definitions:

- To shorten individual specification statements. For example, when argument lists to functions get too long (say, greater than four), then it probably means the function being defined is “doing too much.”
- To “chunkify” and enable reuse of concepts. When a long expression (say, involving more than two logical operators and three function symbols) appears multiple (say, more than two) times, then it probably means that chunk of information can be given a name and the name reused accordingly.
- To postpone specifying certain details. When you find yourself going into too much depth while specifying one component of the system in neglect of specifying the rest of the system, then introduce a placeholder term to be defined later.

Notation

If your primary purpose in specifying is the tangible end-product, i.e., the specification, and you have chosen a particular formal method to use, stick to its notation. Presumably you chose this formal method for its brand of expressiveness or for its known applicability to your problem domain. A carefully designed specification language should have just the right number and kinds of syntactic constructs to let you express all of what you want to say. The constructs provided by the language let you highlight those aspects of the system that are important to record, e.g., side effects in a Larch **modifies** clause. At the same time, they also force you to express yourself in a stylized way using a restricted vocabulary. So, once in a while the notation may force you to express something more awkwardly or more verbosely than you wish; however, either situation may actually be a sign to rethink your abstractions and decompositions.

If you are primarily interested in gaining a deeper understanding of your system through the use of formalism and the formal specification you write is a means toward your end, then do not feel overly constrained by notation. You might happen, not necessarily out of choice, to be using a formal method not specifically designed for your problem domain. If there is a concept you want to express and you cannot express it easily in the given notation, invent some convenient syntax, say what you want, and defer giving it a formal meaning till later. Don’t let notation get in the way of your making progress in writing your specification. On the other hand, don’t forget to define your inventions. It may be at odds with the rest of the semantics. (If you’re lucky, however, you will have thought of a new specification language idiom that is more generally useful than for just your problem at hand.)

Since no one method is suitable for specifying all aspects of a system or all kinds of systems, you might choose to resort to the only practical strategy known today: to mix methods, and hence, to use a mix of notations. For example, you might use Z to specify the static properties of your system (state space); and CSP, its dynamic behavior (sequences of state transitions). Mixing methods,




Figure 1: A Commuting Diagram

however, is dangerous: If you use different notations from different methods, then because they are based on different semantics (e.g., state machines and process algebras), you are less likely to detect that you have specified something that is actually semantically inconsistent. Combining different formal methods is a subject of current research (e.g., see [1, 16, 20]).

Proofs

Most likely you will not be proving theorems about your system from your specifications, but if you are, the first difficult aspect about doing proofs is knowing how formal to be. For realistic systems or large examples, it's impractical to do a completely formal proof, in the strictest sense of “formal” as used in mathematical logic. What you should strive for when writing out an informal proof is to justify each proof step that in principle could be formalized.

Given that you are doing only informal proofs, the second difficult aspect is knowing when you can skip steps. Some steps are “obvious” but others are not. Also, what may seem “obvious” often reflects a hole in your argument.

It is possible to do large formal proofs using machine aids like proof checkers and theorem provers. There is of course a tradeoff between the effort needed to learn to use one of these tools and its input language and underlying logic and the benefit gained by doing the more formal proof. If what you are trying to prove is critical, it may pay to invest the time and energy; moreover, this cost need be paid only once, the first time. If you plan to do more than one (critical) proof, it may be worth your while. Finally, using machine aids keeps you honest because they do not let you skip steps.

Choosing the degree of formality and how much proof detail to give takes experience and practice, gained by both reading other people's proofs and constructing your own. A background in mathematics usually helps.

There are common proof techniques that you should have in your arsenal: proof by induction, case analysis, proof by contradiction, and equational reasoning (substituting equals for equals). You should be familiar with natural deduction though you probably would use it for only small, local proofs.

Finally, the familiar commuting diagram from mathematics plays a central role in proofs of correctness for software systems. For example, an interpretation for Fig. 1 in the context of state machines is to suppose that f is an action of a concrete machine on the concrete state x . If $\langle x, f, x' \rangle$ is a state transition of the concrete machine, then there exists an abstract action g such that $\langle A(x), g, A(x') \rangle$ is a state transition of the abstract machine.

In the context of abstract data types, the interpretation is that given that x is a concrete representation for y , the concrete function f *implements* the abstract function g under the abstraction function A . That is,

$$A(f(x)) = g(A(x))$$

More elaborate diagrams, for example, that allow sequences of actions rather than single actions generalize this basic idea. The “CLInc Stack” case study [4] of proving the correctness of the implementation of a small programming language down to the hardware level relies fundamentally on a stack of commuting diagrams.

5.2. The Details

I now turn to the nitty gritty of specification: getting the technical details right.

Logical Errors

Common logical errors that I have seen specifiers (including myself) make involve implication and quantification.

Implication. Remember that *false* implies anything so that

$$false \Rightarrow \dots$$

is vacuously true, and that anything implies *true* so that

$$\dots \Rightarrow true$$

reduces to true.

Quantifiers. Problem spots include nested quantifiers, ordering of quantifiers (especially modal operators for a temporal logic), and combining quantifiers and implication (e.g., what happens to a formula when bringing a quantifier outside an implication). Another confusion arises when qualifying a quantified variable with set membership, \in . That is,

$$\forall x \in T . P(x)$$

translates to

$$\forall x . x \in T \Rightarrow P(x)$$

but

$$\exists x \in T . P(x)$$

translates to

$$\exists x . x \in T \wedge P(x)$$

If you have a complicated predicate with a lot of embedded quantifiers, you may find it helpful to break the predicate into pieces, where each piece is in prenex normal form and has only one or two quantifiers.

Properties of sets, functions, and relations

When specifying objects such as sets, bags, and sequences that are collections of objects you may be prone to making the following common errors.

Saying

$$x \in s'$$

in the post-condition of an *insert* operation on sets is not enough. It does not say that elements in the set that were originally in s are still there.

Suppose you are specifying the behavior of a *remove* operation, which extracts and returns an element, x , from a set, s : Saying

$$s' = s - \{x\}$$

in the post-condition is too weak. You need to say

$$s' = s - \{x\} \wedge x \in s$$

since in the first case x may not be a member of s and the post-condition could hold by returning an arbitrary value; in that case, the set would also not change in value, probably not the intended behavior for a *remove* operation.

Saying

$$s - s' = \{x\}$$

is also not strong enough. Here you need to add that s' is a proper subset of s :

$$s - s' = \{x\} \wedge s' \subset s$$

since the first case allows s' to have extra elements.

Some specification languages allow functional notation to be used for relations that are not functions. If you try to do ordinary mathematical (and specifically algebraic) reasoning with formulae written using that notation, you are headed for trouble. For example, suppose *choose* is a relation that is not a function; it returns some element from a set. Saying something like

$$f(\text{choose}(s)) \wedge g(\text{choose}(s))$$

in the post-condition of an operation is weaker than saying

$$\exists x . x = \text{choose}(s) \wedge f(x) \wedge g(x)$$

since in the first case the different occurrences of *choose* could return different values. Of course if *choose* is a function, then it is guaranteed to return the same value.

Recursive definitions, commonly found in algebraic specifications, may at first look puzzling. For example, in specifying the *delete* operation for sets

$$\begin{aligned} \text{delete}(\text{insert}(s, e_1), e_2) = & \text{if } e_1 = e_2 \text{ then } \text{delete}(s, e_2) \\ & \text{else } \text{insert}(\text{delete}(s, e_2), e_1) \end{aligned}$$

a common error is to forget to reapply *delete* recursively if e_1 and e_2 are equal or to forget to “reinsert” e_1 if they are not. Without reapplying the *delete* you get a bag, not a set, and without reinserting the e_1 you lose an element from the set.

6. Summary

The process of writing specifications borrows from and is similar to the processes of writing prose, writing programs, and writing mathematics. You need to worry about the big picture (e.g., the overall structure, organization, and meaning of concepts) as well as the fine details (e.g., punctuation, spelling, and special symbols). There are rules that you must always obey and rules that you may break once in a while. There are stylistic conventions to learn and follow. As with writing prose, programs, and mathematics, writing specifications well takes practice and patience.

There are many books on how to write good prose (e.g., [19]) and even some on how to write good programs (e.g., [3]). This paper is my attempt, and perhaps the first such attempt in the formal methods community, to cull out some common rules-of-thumb for writing specifications. Maybe these hints can serve eventually as a basis for a set of organized guidelines for specifiers.

Acknowledgments

My views on writing specifications have been greatly influenced by John Guttag and Jim Horning, originating from our joint paper written in 1982 [7]. More recently, David Garlan and Daniel Jackson have helped identify common specification stumbling blocks that we see our students frequently face in our teaching of the Models, Methods, and Analysis courses for the Carnegie Mellon MSE program. I thank David and Daniel also for their useful feedback on an earlier draft of this paper.

I thank Jim Horning, Leslie Lamport, David Parnas, and Leo Marcus for their critical remarks on the AMAST’95 version of this paper. Jim, as usual, helped polish my writing. They all helped

clarify my comments on nondeterminism and the Larch **modifies** clause. They also convinced me that the AMAST'95 paper's original title was terribly misleading.

Finally, I thank all the students with whom I have worked over the years for their interest and patience in trying their hand at formal specification.

References

- [1] DIS 8807. Information systems processing—open systems interconnection—lotos. Technical report, International Standards Organization, 1987.
- [2] H. G. Baker. List Processing in Real Time on a Serial Computer. *Communications of the ACM*, 21(4):280–294, 1978.
- [3] J. Bentley. *Programming Pearls*. Addison-Wesley, 1986.
- [4] W.R. Bevier, W.A. Hunt, Jr., J S. Moore, and W.D. Young. An approach to systems verification. *Journal of Automated Reasoning*, 5:411–428, 1989.
- [5] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, February 1990.
- [6] D. Garlan. Preconditions for understanding. In *Proceedings of the Sixth Int'l Conf. on Software Specification and Design*, pages 242–245, October 1991.
- [7] J.V. Guttag, J.J. Horning, and J.M. Wing. Some remarks on putting formal specifications to productive use. *Science of Computer Programming*, 2(1):53–68, October 1982.
- [8] C.A.R. Hoare. *Communicating Sequential Processes*. Prentice-Hall International, 1985.
- [9] C.B. Jones. *Software Development: A Rigorous Approach*. Prentice-Hall International, 1980.
- [10] C.B. Jones. *Systematic Software Development Using VDM*, Chapter 15. Prentice-Hall International, 1986.
- [11] B. Liskov and J. Guttag. *Abstraction and Specification in Program Development*. McGraw-Hill/MIT Press, 1986.
- [12] N. Lynch and M. Tuttle. Hierarchical correctness proofs for distributed algorithms. Technical report, MIT Laboratory for Computer Science, Cambridge, MA, April 1987.
- [13] A.J.R.G. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer-Verlag, 1980.
- [14] L. Mummert, J.M. Wing, and M. Satyanarayanan. Using belief to reason about cache coherence. In *Proceedings of the Symposium on Principles of Distributed Computing*, pages 71–80, August 1994. Also CMU-CS-94-151, May 1994.
- [15] Scott M. Nettles and James W. O'Toole. Real-Time Replication Garbage Collection. In *SIGPLAN Symposium on Programming Language Design and Implementation*. ACM, June 1993.

- [16] M. Nielsen, K. Havelund, K.R. Wagner, and C. George. The RAISE language, method and tools. *Formal Aspects of Computing*, 1:85–114, 1989.
- [17] John Scheid and Steven Holtsberg. Ina Jo specification language reference manual. Technical Report TM-6021/001/06, Paramax Systems Corporation, A Unisys Company, June 1992.
- [18] J.M. Spivey. *Introducing Z: a Specification Language and its Formal Semantics*. Cambridge University Press, 1988.
- [19] W. Strunk and E.B.White. *The Elements of Style*. The Macmillan Company, 1972. second edition.
- [20] P. Zave and M. Jackson. Conjunction as composition. *ACM Trans. on Soft. Engineering and Methodology*, 2(4):379–411, October 1993.