

Design by Contract: building reliable software

E quipped with the basic concepts of class, object and genericity, you can by now write software modules that implement possibly parameterized types of data structures. Congratulations. This is a significant step in the quest for better software architectures.

But the techniques seen so far are not sufficient to implement the comprehensive view of quality introduced at the beginning of this book. The quality factors on which we have concentrated — reusability, extensibility, compatibility — must not be attained at the expense of reliability (*correctness* and *robustness*). Although, as recalled next, the reliability concern was visible in many aspects of the discussion, we need more.

The need to pay more attention to the semantic properties of our classes will be particularly clear if you remember how classes were defined: as implementations of abstract data types. The classes seen so far consist of attributes and routines, which indeed represent the functions of an ADT specification. But an ADT is more than just a list of available operations: remember the role played by the semantic properties, as expressed by the axioms and preconditions. They are essential to capture the true nature of the type's instances. In studying classes, we have — temporarily — lost sight of this semantic aspect of the ADT concept. We will need to bring it back into the method if we want our software to be not just flexible and reusable, but also correct and robust.

Assertions and the associated concepts, explained in this chapter, provide some of the answer. Although not foolproof, the mechanisms presented below provide the programmer with essential tools for expressing and validating correctness arguments. The key concept will be **Design by Contract**: viewing the relationship between a class and its clients as a formal agreement, expressing each party's rights and obligations. Only through such a precise definition of every module's claims and responsibilities can we hope to attain a significant degree of trust in large software systems.

In reviewing these concepts, we shall also encounter a key problem of software engineering: how to deal with run-time errors — with contract violations. This leads to the subject of *exception handling*, covered in the next chapter. The distribution of roles between the two chapters roughly reflects the distinction between the two components of reliability: as you will recall, correctness was defined as the software's ability to perform according to its specification, and robustness as its ability to react to cases not included in the specification. Assertions (this chapter) generally cover correctness, and exceptions (next chapter) generally cover robustness.

Some important extensions to the basic ideas of Design by Contract will have to wait until the presentation of inheritance, polymorphism and dynamic binding, enabling us to go from contracts to *subcontracting*.

11.1 BASIC RELIABILITY MECHANISMS

The preceding chapters already introduced a set of techniques that directly address the goal of producing reliable software. Let us review them briefly: it would be useless to consider more advanced concepts until we have put in place all the basic reliability mechanisms.

First, the defining property of object technology is an almost obsessive concern with the *structure* of software systems. By defining simple, modular, extensible architectures, we make it easier to ensure reliability than with conformed structures as often result from earlier methods. In particular the effort to limit inter-module communication to the strict minimum was central to the discussion of modularity that got us started; it resulted in the prohibition of such common reliability risks as global variables, and in the definition of restricted communication mechanisms, the client and inheritance relations. The general observation is that the single biggest enemy of reliability (and perhaps of software quality in general) is complexity. Keeping our structures as simple as possible is not enough to ensure reliability, but it is a necessary condition. So the discussion of the previous chapters provides the right starting point for the systematic effort of the present one.

Also necessary if not sufficient is the constant emphasis on making our software *elegant* and *readable*. Software texts are not just written, they are read and rewritten many times; clarity and simplicity of notation, such as have been attempted in the language constructs introduced so far, are a required basis for any more sophisticated approach to reliability.

Another indispensable weapon is automatic memory management, specifically *garbage collection*. The chapter on memory management explained in detail why, for any system that creates and manipulates dynamic data structures, it would be dangerous to rely on manual reclamation (or no reclamation). Garbage collection is not a luxury; it is a crucial reliability-enhancing component of any O-O environment.

The same can be said of another technique presented (in connection with genericity) in the last chapter: static typing. Without statically enforced type rules, we would be at the mercy of run-time typing errors.

All these techniques provide the necessary basis, from which we can now take a closer look at what it will take for a software system to be correct and robust.

11.2 ABOUT SOFTWARE CORRECTNESS

We should first ask ourselves what it *means* for a software element to be correct. The observations and deductions that will help answer this question will seem rather trivial at first; but let us not forget the comment (made once by a very famous scientist) that scientific reasoning is nothing but the result of starting from ordinary observations and continuing with simple deductions — only very patiently and stubbornly.

Assume someone comes to you with a 300,000-line C program and asks you “Is this program correct?”. There is not much you can answer. (If you are a consultant, though, try answering “no” and charging a high fee. You might just be right.)

To consider the question meaningful, you would need to get not only the program but also a precise description of what it is supposed to do — a *specification*.

The same comment is applicable, of course, regardless of the size of a program. The instruction $x := y + 1$ is neither correct nor incorrect; these notions only make sense with respect to a statement of what one expects from the instruction — what effect it is intended to have on the state of the program variables. The instruction is correct for the specification

“*Make sure that x and y have different values*”

but it is incorrect vis-à-vis the specification

“*Make sure that x has a negative value*”

(since, assuming that the entities involved are integers, x may end up being non-negative after the assignment, depending on the value of y).

These examples illustrate the property that must serve as the starting point of any discussion of correctness.

Software Correctness property

Correctness is a relative notion.

A software system or software element is neither correct nor incorrect per se; it is correct or incorrect with respect to a certain specification. Strictly speaking, we should not discuss whether software elements are *correct*, but whether they are *consistent* with their specifications. This discussion will continue to use the well-accepted term “correctness”, but we should always remember that the question of correctness does not apply to software elements; it applies to pairs made of a software element and a specification.

In this chapter we will learn how to express such specifications through **assertions**, to help us assess the correctness of our software. But we will go further. It turns out (and only someone who has not practiced the approach will think of this as a paradox) that just writing the specification is a precious first step towards *ensuring* that the software actually meets it. So we will derive tremendous benefits from writing the assertions at the same time as we write the software — or indeed before we write the software. Among the consequences we will find the following:

- Producing software that is correct from the start because it is designed to be correct. [Mills 1975]. The title of an article written by Harlan D. Mills (one of the originators of “Structured Programming”) in the nineteen-seventies provides the right mood: *How to write correct programs and know it*. To “know it” means to equip the software, at the time you write it, with the arguments showing its correctness.
- Getting a much better understanding of the problem and its eventual solutions.
- Facilitating the task of software documentation. As we will see later in this chapter, assertions will play a central part in the object-oriented approach to documentation.
- Providing a basis for systematic testing and debugging.

The rest of this chapter explores these applications.

A word of warning: C, C++ and some other languages (following the lead of Algol W) have an “assert” instruction that tests whether a certain condition holds at a certain stage of the software’s execution, and stops execution if it does not. Although relevant to the present discussion, this concept represents only a small part of the use of assertions in the object-oriented method. So if like many other software developers you are familiar with such instructions but have not been exposed to the more general picture, almost all the concepts of this chapter will be new.

11.3 EXPRESSING A SPECIFICATION

We can turn the preceding informal observations into a simple mathematical notation, borrowed from the theory of formal program validation, and precious for reasoning about the correctness of software elements.

Correctness formulae

Let A be some operation (for example an instruction or a routine body). A **correctness formula** is an expression of the form

$\{P\} A \{Q\}$

denoting the following property, which may or may not hold:

Meaning of a correctness formula $\{P\} A \{Q\}$

“Any execution of A , starting in a state where P holds, will terminate in a state where Q holds.”

Correctness formulae (also called *Hoare triples*) are a mathematical notation, not a programming construct; they are not part of our software language, but only designed to guide us through this discussion by helping to express properties of software elements.

In $\{P\} A \{Q\}$ we have seen that A denotes an operation; P and Q are properties of the various entities involved, also called assertions (the word will be defined more precisely later). Of the two assertions, P is called the precondition and Q the postcondition. Here is a trivial correctness formula (which, assuming that x is an integer entity, holds):

$$\{x \geq 9\} \ x := x + 5 \ \{x \geq 13\}$$

The use of correctness formulae is a direct application of the Software Correctness Property. What the Property stated informally — that correctness is only meaningful relative to a particular specification — correctness formulae turn into a form that is directly usable for working on the software: from now on the discourse about software correctness will not be about individual software elements A , but about triples containing a software element A , a precondition P and a postcondition Q . The sole aim of the game is to establish that the resulting $\{P\} A \{Q\}$ correctness formulae hold.

The number 13 appearing in the postcondition is not a typo! Assuming a correct implementation of integer arithmetic, the above formula holds: if $x \geq 9$ is true before the instruction, $x \geq 13$ will be true after the instruction. Of course we can assert more interesting things with the given precondition, the most interesting postcondition is the strongest possible one, here $x \geq 14$; with the given postcondition, the most interesting precondition is the weakest possible one, here $x \geq 8$. From a formula that holds, you can always get another one by strengthening the precondition or weakening the postcondition. We will now examine more carefully these notions of “stronger” and “weaker”.

Weak and strong conditions

One way to look at a specification of the form $\{P\} A \{Q\}$ is to view it as a job description for A — an ad in the paper, which states “We are looking for someone whose work will be to start from initial situations as characterized by P , and deliver results as defined by Q ”.

Here is a small quiz to help you sharpen your understanding of the concepts.

Assume one of your friends is looking for a job and comes across several such ads, all with similar salary and benefits, but differing by their P s and Q s. (Tough times have encouraged the companies that publish the ads to resort to this notation, which they like for its mathematical compactness since the newspaper charges by the word.) Like everyone else, your friend is lazy, that is to say, wants to have the easiest possible job. He is asking for your advice, always a dangerous situation. What should you recommend for P : choose a job with a weak precondition, or a strong one? Same question for the postcondition Q . (The answers appear right after this, but do take the time to decide the issue for yourself before turning the page.)

The precondition first. From the viewpoint of the prospective employee — the person who has to perform what has been called A — the precondition P defines the conditions under which the required job will start or, to put it differently, the set of cases that have to be handled. So a strong P is good news: it means that you only have to deal with a limited set of situations. The stronger the P , the easier for the employee. In fact, the perfect sinecure is the job defined by

Sinecure 1
 $\{False\} A \{...\}$

The postcondition has been left unspecified because it does not matter what it is. Indeed if you ever see such an ad, do not even bother reading the postcondition: *take the job right away*. The precondition *False* is the strongest possible assertion, since it is never satisfied in any state. Any request to execute A will be incorrect, and the fault lies not with the agent responsible for A but with the requester — the client — since it did not observe the required precondition, for the good reason that it is impossible to observe it. Whatever A does or does not do may be *useless*, but is always *correct* — in the sense, defined earlier, of being consistent with the specification.

The above job specification is probably what a famous police chief of a Southern US city had in mind, a long time ago, when, asked by an interviewer why he had chosen this career, he replied: “Obvious — it is the only job where the customer is always wrong”.

For the postcondition Q , the situation is reversed. A strong postcondition is bad news: it indicates that you have to deliver more results. The weaker the Q , the better for the employee. In fact, the second best sinecure in the world is the job defined, regardless of the precondition, by

Sinecure 2
 $\{...\} A \{True\}$

The postcondition *True* is the weakest possible assertion, satisfied by all states.

The notions of “stronger” and “weaker” are formally defined from logic: P_1 is said to be stronger than P_2 , and P_2 weaker than P_1 , if P_1 implies P_2 and they are not equal. As every proposition implies *True*, and *False* implies every proposition, it is indeed legitimate to speak of *True* as the weakest and *False* as the strongest of all possible assertions.

Why, by the way, is Sinecure 2 only the “second best” job in the world? The reason has to do with a fine point that you may have noticed in the definition of the meaning of $\{P\} A \{Q\}$ on the preceding page: termination. The definition stated that the execution must *terminate* in a state satisfying Q whenever it is started in a state satisfying P . With Sinecure 1 there are no states satisfying P , so it does not matter what A does, even if it is a program text whose execution would go into an infinite loop or crash the computer. Any A will be “correct” with respect to the given specification. With Sinecure 2, however, there

must be a final state; that state does not need to satisfy any specific properties, but it must exist. From the viewpoint of whoever has to perform *A*: you need to do nothing, *but you must do it in finite time*.

Readers familiar with theoretical computing science or program proving techniques will have noted that the *(P) A (Q)* notation as used here denotes **total correctness**, which includes termination as well as conformance to specification. (The property that a program will satisfy its specification if it terminates is known as partial correctness.) See [M 1990] for a detailed presentation of these concepts.

The discussion of whether a stronger or weaker assertion is “bad news” or “good news” has taken the viewpoint of the prospective employee. If, changing sides, we start looking at the situation as if we were the employer, everything is reversed: a weaker precondition will be good news, as it means a job that handles a broader set of input cases; so will be a stronger postcondition, as it means more significant results. This reversal of criteria is typical of discussions of software correctness, and will reappear as the central notion of this chapter: *contracts* between client and supplier modules, in which a benefit for one is an obligation for the other. To produce effective and reliable software is to draw up the contract representing the best possible compromise in all applicable client-supplier communications.

11.4 INTRODUCING ASSERTIONS INTO SOFTWARE TEXTS

Once we have defined the correctness of a software element as the consistency of its implementation with its specification, we should take steps to include the specification, together with the implementation, in the software itself. For most of the software community this is still a novel idea: we are accustomed to programs as defining the operations that we command our hardware-software machines to execute for us (the *how*); it is less common to treat the description of the software’s purposes (the *what*) as being part of the software itself.

To express the specification, we will rely on assertions. An assertion is an expression involving some entities of the software, and stating a property that these entities may satisfy at certain stages of software execution. A typical assertion might express that a certain integer has a positive value or that a certain reference is not void.

Mathematically, the closest notion is that of predicate, although the assertion language that we shall use has only part of the power of full predicate calculus.

Syntactically, the assertions of our notation will simply be boolean expressions, with a few extensions. One of these extensions, the *old* notation, is introduced later in this chapter. Another is the use of the semicolon, as in

n > 0 ; x /= Void

The meaning of the semicolon is equivalent to that of an *and*. As between declarations and instructions, the semicolon is actually optional, and we will omit it when assertion clauses appear on separate lines; just consider that there is an implicit *and* between successive assertion lines. These conventions facilitate identification of the individual components of an assertion. It is indeed possible, and usually desirable, to label these components individually, as in

Positive: n > 0
Not void: x /= Void

If present, the labels (such as *Positive* and *Not void* in this example) will play a role in the run-time effect of assertions — to be discussed later in this chapter — but for the moment they are mainly there for clarity and documentation.

The next few sections will review this principal application of assertions: as a conceptual tool enabling software developers to construct correct systems and to document why they are correct.

11.5 PRECONDITIONS AND POSTCONDITIONS

The first use of assertions is the semantic specification of routines. A routine is not just a piece of code; as the implementation of some function from an abstract data type specification, it should perform a useful task. It is necessary to express this task precisely, both as an aid in designing it (you cannot hope to ensure that a routine is correct unless you have specified what it is supposed to do) and, later, as an aid to understanding its text.

You may specify the task performed by a routine by two assertions associated with the routine: a *precondition* and a *postcondition*. The precondition states the properties that must hold whenever the routine is called; the postcondition states the properties that the routine guarantees when it returns.

A stack class

An example will enable us to become familiar with the practical use of assertions. In the previous chapter, we saw the outline of a generic stack class, under the form

class STACK [G] feature
... Declaration of the features:
count, empty, full, put, remove, item
end

An implementation will appear below. Before considering implementation issues, however, it is important to note that the routines are characterized by strong semantic properties, independent of any specific representation. For example:

- Routines *remove* and *item* are only applicable if the number of elements is not zero.
- *put* increases the number of elements by one; *remove* decreases it by one.

Such properties are part of the abstract data type specification, and even people who do not use any approach remotely as formal as ADTs understand them implicitly. But in common approaches to software construction software texts reveal no trace of them. Through routine preconditions and postconditions you can turn them into explicit elements of the software.

We will express preconditions and postconditions as clauses of routine declarations introduced by the keywords *require* and *ensure* respectively. For the stack class, leaving the routine implementations blank for the time being, this gives:

```
indexing
  description: "Stacks: Dispenser structures with a Last-In, First-Out %
    %access policy"
class STACK1 [G] feature -- Access
  count: INTEGER
    -- Number of stack elements

  item: G is
    -- Top element
    require
      not empty
    do
      ...
    end
  feature -- Status report
    empty: BOOLEAN is
      -- Is stack empty?
      do ... end

  full: BOOLEAN is
    -- Is stack representation full?
    do
      ...
    end
  feature -- Element change
    put (x: G) is
      -- Add x on top.
      require
        not full
      do
        ...
      ensure
        not empty
        item = x
      end
    end
  ensure
    count = old count + 1
  end

  remove is
    -- Remove top element.
    require
      not empty
    do
      ...
    ensure
      not full
      count = old count - 1
    end
  end
end
```

Both the **require** and the **ensure** clauses are optional; when present, they appear at the places shown. The **require** appears before the **local** clause, if present. The next sections explain in more detail the meaning of preconditions and postconditions.

Note the division into several **feature** clauses, useful to group the features into categories indicated by the clauses' header comments. **Access**, **Status report** and **Element change** are some of a dozen or so standard categories used throughout the libraries and, whenever applicable, subsequent examples in this book.

More on **feature** categories in "A stack class", page 348.

Preconditions

A precondition expresses the constraints under which a routine will function properly. Here:

- **put** may not be called if the stack representation is full.
- **remove** and **item** may not be applied to an empty stack.

A precondition applies to all calls of the routine, both from within the class and from clients. A correct system will never execute a call in a state that does not satisfy the precondition of the called routine.

Postconditions

A postcondition expresses properties of the state resulting from a routine's execution. Here:

- After a **put**, the stack may not be empty, its top is the element just pushed, and its number of elements has been increased by one.
- After a **remove**, the stack may not be full, and its number of elements has been decreased by one.

The presence of a postcondition clause in a routine expresses a guarantee on the part of the routine's implementor that the routine will yield a state satisfying certain properties, assuming it has been called with the precondition satisfied.

A special notation, **old**, is available in postconditions; **put** and **remove** use it to express the changes to **count**. The notation **old *e***, where *e* is an expression (in most practical cases an attribute), denotes the value that *e* had on routine entry. Any occurrence of *e* not preceded by **old** in the postcondition denotes the value of the expression on exit. The postcondition of **put** includes the clause

count = old count + 1

to state that **put**, when applied to any object, must increase by one the value of the **count** field of that object.

A pedagogical note

If you are like most software professionals who get exposed to these ideas for the first time, you may be itching to know what effect, if any, the assertions have on the execution of the software, and in particular what happens if one of them gets violated at run time — if **full** is true when someone calls **put**, or **empty** is true when **put** terminates one of its executions. It is too early to give the full answer but as a preview we can use the lawyer's favorite: *it depends*.

More precisely, it depends on what you want. You may decide to treat assertions purely as comments, with no effect on the software's execution; then a run-time assertion violation will remain undetected. But it is also possible to use assertions to check that everything goes according to plan: then during execution the environment will automatically monitor that all assertions held when they should, and if one does not it will trigger an exception, usually terminating execution and printing a message indicating clearly what happened. (It is also possible to include an exception handling clause that will try to recover from the exception and continue execution; exception handling is discussed in detail in the next chapter.) To specify the policy that you want — no assertion checking, or assertion monitoring at one of various possible levels — you will use a compilation option, which you can set separately for each class.

The full details of run-time assertion monitoring do appear later in this chapter. But it would be a mistake to attach too much importance to this aspect at this stage (one of the reasons why you were warned earlier not to think too much about the C notion of assertion if that has been your only exposure to the concept). Other aspects of assertions demand our attention first. We have only started to see assertions as a technique to help us get our software right in the first place; we still have much to discover of their *methodological* role as built-in guardians of reliability. The question of what happens if we do fail (in particular if an assertion, in spite of all our efforts, is not satisfied at some execution instant) is important too, but only after we have done all we could to prevent it from arising.

So (although it is never bad to think ahead) you do not need at this point to be too preoccupied by such questions as the possible performance penalty implied by the **old** construct. Must the run-time system preserve values before we start a routine, just to be able to evaluate an **old** expression appearing in the postcondition? *It depends*: in some circumstances (for example testing and debugging) it will indeed be useful to evaluate assertions; in others (for example production runs of fully validated systems) you can treat them as mere annotations to the software text.

All that counts for the next few sections is the methodological contribution of assertions, and of the associated method of Design by Contract: as a conceptual tool for analysis, design, implementation and documentation, helping us to build software in which **reliability** is **built-in**, rather than achieved or attempted after the fact through debugging; in Mill's's terms, enabling us to build correct programs and know it.

11.6 CONTRACTING FOR SOFTWARE RELIABILITY

Defining a precondition and a postcondition for a routine is a way to define a *contract* that binds the routine and its callers.

Rights and obligations

By associating clauses **require** *pre* and **ensure** *post* with a routine *r*, the class tells its clients:

"If you promise to call *r* with *pre* satisfied then I, in return, promise to deliver a final state in which *post* is satisfied."

In relations between people or companies, a contract is a written document that serves to clarify the terms of a relationship. It is really surprising that in software, where precision is so important and ambiguity so risky, this idea has taken so long to impose itself. A precondition-postcondition pair for a routine will describe the contract that the routine (the *supplier* of a certain service) defines for its callers (the *clients* of that service).

Perhaps the most distinctive feature of contracts as they occur in human affairs is that any good contract entails obligations as well as benefits for both parties — with an obligation for one usually turning into a benefit for the other. This is true of contracts between classes, too:

- The precondition binds the client: it defines the conditions under which a call to the routine is legitimate. It is an *obligation* for the client and a *benefit* for the supplier.
- The postcondition binds the class: it defines the conditions that must be ensured by the routine on return. It is a benefit for the client and an obligation for the supplier.

The benefits are, for the client, the guarantee that certain properties will hold after the call; for the supplier, the guarantee that certain assumptions will be satisfied whenever the routine is called. The obligations are, for the client, to satisfy the requirements as stated by the precondition; for the supplier, to do the job as stated by the postcondition.

Here is the contract for one of the routines in our example:

<i>put</i>		OBLIGATIONS	BENEFITS
<i>Client</i>	(<i>Satisfy precondition:</i>) Only call <i>put</i> (<i>x</i>) on a non-full stack.	(<i>From postcondition:</i>) Get stack updated: not empty, <i>x</i> on top (<i>item</i> yields <i>x</i> , <i>count</i> increased by 1).	
<i>Supplier</i>	(<i>Satisfy postcondition:</i>) Update stack representation to have <i>x</i> on top (<i>item</i> yields <i>x</i>), <i>count</i> increased by 1, not empty.		(<i>From precondition:</i>) Simpler processing thanks to the assumption that stack is not full.

A routine contract:
routine *put* for a stack class

Zen and the art of software reliability: guaranteeing more by checking less

Although you may not have noticed it yet, one of the contract rules given goes against the generally accepted wisdom in software engineering: shocking at first to many, it is among the method's main contributions to software reliability and deserves emphasis.

The rule reflects the above observation that the precondition is a *benefit* for the supplier and is expressed in the bottom-right box of the table: if the client's part of the

contract is not fulfilled, that is to say if the call does not satisfy the precondition, then the class is not bound by the postcondition. In this case the routine may do what it pleases: return any value; loop indefinitely without returning a value; or even crash the execution in some wild way. This is the case in which (in reference to the discussion at the beginning of this chapter) “the customer is wrong”.

The first advantage of this convention is that it considerably simplifies the programming style. Having specified as a precondition the constraints which calls to a routine must observe, you, the class developer, may assume when writing the routine body that the constraints are satisfied; you do not need to test for them in the body. So if a square root function, meant to produce a real number as a result, is of the form

```
sqr(x: REAL): REAL is
-- Square root of x
require
  x >= 0
do ... end
```

you may write the algorithm for computing the square root without any concern for the case in which x is negative; this is taken care of by the precondition and becomes the responsibility of your clients. (At first sight this may appear dangerous; but read on.)

Actually the method of Design by Contract goes further. Writing the **do** clause of the routine under the form

```
if x < 0 then
  “Handle the error, somehow”
else
  “Proceed with normal square root computation”
end
```

is not just unnecessary but unacceptable. This may be expressed as a methodological rule:

Non-Redundancy principle

Under no circumstances shall the body of a routine ever test for the routine’s precondition.

This rule is the reverse of what many software engineering or programming methodology textbooks advocate, often under the name *defensive programming* — the idea that to obtain reliable software you should design every component of a system so that it protects itself as much as possible. Better check too much, this approach holds, than not enough; one is never too careful when dealing with strangers. A redundant check might not help, but at least it will not hurt.

Design by Contract follows from the opposite observation: redundant checks can and indeed will hurt. Of course this will at first seem strange; the natural reaction is to think that an extra check — for example routine *sqr* containing the above conditional instruction testing for $x < 0$ even though callers have been instructed to ensure $x \geq 0$ —

may at worst be useless, but cannot possibly cause any damage. Such a comment, however, comes from a microscopic understanding of reliability, focused on individual software elements such as the *sqr* routine. If we restrict our view to the narrow world of *sqr*, then the routine seems more robust with the extra test than without it. But the world of a system is not restricted to a routine; it contains a multitude of routines in a multitude of classes. To obtain reliable systems we must go from the microscopic view to a macroscopic view encompassing the entire architecture.

If we take this global view, *simplicity* becomes a crucial criterion. As was noted at the beginning of this chapter, complexity is the major enemy of quality. When we bring in this concern, possibly redundant checks do not appear so harmless any more! Extrapolated to the thousands of routines of a medium-size system (or the tens or hundreds of thousands of routines of a larger one), the **if $x < 0$ then ... of *sqr***, innocuous at first sight, begins to look like a monster of useless complexity. By adding possibly redundant checks, you add more software; more software means more complexity, and in particular more sources of conditions that could go wrong; hence the need for more checks, meaning more software; and so on ad infinitum. If we start on this road only one thing is certain: we will never obtain reliability. The more we write, the more we will have to write.

To avoid this infinite chase we should never start it. With Design by Contract you are invited to identify the consistency conditions that are necessary to the proper functioning of each client-supplier cooperation (each contract); and to specify, for each one of these conditions, **whose responsibility it is** to enforce it: the client’s, or the supplier’s. The answer may vary, and is partly a matter of design style: advice will be given below on how best to choose it. But once you have made the decision, you should stick to it: if a correctness requirement appears in the precondition, indicating that the requirement is part of the client’s responsibility, there must not be a corresponding test in the routine; and if it is not in the precondition, then the routine must check for the requirement.

Defensive programming appears in contrast to cover up for the lack of a systematic approach by blindly putting in as many checks as possible, furthering the problem of reliability rather than addressing it seriously.

Redundant checking, it should be noted, is a standard technique in hardware. The difference is that in a hardware system some object that was found to be in a correct state at some point may later have its integrity destroyed because of reasons beyond the control of the system itself, such as interference from another system, harmful external event, or simply wear and tear. For that reason it is normal practice, for example, to have both the sender and the receiver of an electronic signal check its integrity.

But no such phenomenon occurs in software: if I can prove or check in some way that o is non-negative whenever *sqr*(o) is called, I do not need to insert a check for $x \geq 0$, where x is the corresponding formal argument, in the body of *sqr*. Nothing will happen to o between the time it is “sent” by the caller and the time it is “received” (under the name x) by the routine. Software does not wear out when used for too long; it is not subject to line loss, to interference or to noise.

Also note that in most cases what is called redundant checking in hardware is not really redundant: one actually applies *different* and complementary verifications, such as a parity check and some other test. Even when the checks are the same they are often

applied by different devices, as in the just mentioned case of a sender and receiver that both check a signal, or in a redundant computer system where several computers perform the same computation, with a voting mechanism to resolve discrepancies.

Another drawback of defensive programming is its costs. Redundant checks imply a performance penalty — often enough in practice to make developers wary of defensive programming regardless of what the textbooks say. If they do make the effort to include these checks, removing some of them later to improve performance will be tedious. The techniques of this chapter will also leave room for extra checks, but if you choose to enable them you will rely on the development environment to carry them out for you. To remove them, once the software has been debugged, it suffices to change a compilation option (details soon). The software itself does not contain any redundant elements.

Aside from performance considerations, however, the principal reason to distrust defensive programming is simply our goal of getting the best possible reliability. For a system of any significant size the individual quality of the various elements involved is not enough; what will count most is the guarantee that for every interaction between two elements there is an explicit roster of mutual obligations and benefits — the contract. Hence the Zen-style paradox of our conclusion: that to get *more* reliability the best policy is often to check *less*.

Assertions are not an input checking mechanism

It is useful here to emphasize a few properties of the approach which, although implicit in the preceding discussion, have been shown by experience to require further explanations. The following comments should help address some of the questions that may have been forming in your mind as you were reading about the basic ideas of Design by Contract.


To avoid a common misunderstanding, make sure to note that each of the contracts discussed holds between a routine (the supplier) and another routine (its caller): we are concerned about software-to-software communication, not software-to-human or software-to-outside-world. A precondition will not take care of correcting user input, for example in a *read_integer* routine that expects the interactive user to enter a positive number. Including in the routine a precondition of the form

require
input > 0

would be wishful thinking, not a reliability technique. Here there is no substitute for the usual condition-checking constructs, including the venerable *if ... then ...*; the exception handling mechanism studied in the next chapter may also be helpful.

Assertions do have a role to play in a solution to this problem of input validation. In line with the criterion of Modular Protection, the method encourages validating any objects obtained from the outside world — from sensors, from user input, from a network... — as close to the source of the objects as possible, using “filter” modules if necessary:

“Modular protection”, page 45.



Using filter modules

In obtaining information from the outside (communication paths shown in color) you cannot rely on preconditions. But part of the task of the input modules shown in grey in the middle of the figure is to guarantee that no information is passed further to the right — to the modules responsible for the system’s actual computations — unless it satisfies the conditions required for correct processing. In this approach there will be ample use of assertions in the software-to-software communication paths represented by the black dotted lines on the right. The postconditions achieved by the routines of the input modules will have to match (or exceed, in the sense of “stronger” defined earlier) the preconditions imposed by the processing routines.

The routines of the filter classes may be compared to security officers in, say, a large government laboratory. To meet experts from the laboratory and ask them technical questions, you must submit to screening procedures. But it is not the same person who checks your authorization level and answers the questions. The physicists, once you have been officially brought into their offices, assume you satisfy the preconditions; and you will not get much help from the guards on theoretical physics.

Assertions are not control structures

Another common misunderstanding, related to the one just discussed, is to think of assertions as control structures — as techniques to handle special cases. It should be clear by now that this is not their role. If you want to write a routine *sqrt* that will handle negative arguments a certain way, and non-negative arguments another way, a *require* clause is not what you need. Conditional instructions (*if ... then ... else ...*) and related constructs to deal with various cases (such as Pascal’s *case ... of ...* or the *inspect* instruction of this book’s notation) are perfectly appropriate for such purposes.

Assertions are something else. They express correctness conditions. If *sqrt* has its precondition, a call for which *x < 0* is not a special case: it is a bug, plain and simple.

Assertion Violation rule (1)
A run-time assertion violation is the manifestation of a bug in the software.

“Bug” is not a very scientific word but is clear enough to anyone in software: we will look for more precise terminology in the next section. For the moment we can pursue the assertion violation rule further by noting a consequence of the contract view:

Assertion violation rule (2)

A precondition violation is the manifestation of a bug in the client.
A postcondition violation is the manifestation of a bug in the supplier.

A precondition violation means that the routine's caller, although obligated by the contract to satisfy a certain requirement, did not. This is a bug in the client itself; the routine is not involved. ("The customer is wrong".) An outside observer might of course criticize the contract as too demanding, as with the unsatisfiable *require False* precondition or our fictitious *Secure 1* example ("the customer is *always* wrong"), but this is too late to argue over the contract: it is the contract, and the client did not observe its part of the deal. So if there is a mechanism for monitoring assertions during execution — as will be introduced shortly — and it detects such a precondition violation, the routine should not be executed at all. It has stated the conditions under which it can operate, and these conditions do not hold; trying to execute it would make no sense.

A postcondition violation means that the routine, presumably called under correct conditions, was not able to fulfill its contract. Here too, the distribution of guilt and innocence is clear, although it is the reverse of the previous one: the bug is in the routine; the caller is innocent.

Errors, defects and other creeping creatures

The appearance of the word "bug" in the preceding analysis of assertion violation causes is a good opportunity to clarify the terminology. In Edsger W. Dijkstra's view, using the word "bug" is a lame attempt by software people to blame someone else by implying that mistakes somehow creep into the software from the outside while the developers are looking elsewhere—as if were not the developers who made the mistakes in the first place.

Yet the term enjoys enduring success, if only because it is colorful and readily understood. Like the rest of the software literature, this book uses it freely. But it is appropriate to complement it by more specific (if more stodgy) terms for cases in which we need precise distinctions.

Terms to denote software woes

- An *error* is a wrong decision made during the development of a software system.
- A *defect* is a property of a software system that may cause the system to depart from its intended behavior.
- A *fault* is the event of a software system departing from its intended behavior during one of its executions.

The causal relation is clear: faults are due to defects, which result from errors.

"Bug" usually has the meaning of defect ("are you sure there remains no other bug in this routine?"). This is the interpretation in this book. But in informal discussions it is also used in the sense of fault ("We have had bug-free operation for the last three weeks") or error ("the bug was that I used an unsorted list").


11.7 WORKING WITH ASSERTIONS

Let us now probe further the use of preconditions and postconditions, continuing with fairly elementary examples. Assertions, some simple, some elaborate, will be pervasive in the examples of the following chapters.

A stack class

The assertion-equipped *STACK* class was left in a sketchy form (*STACK1*). We can now come up with a full version including a spelled out implementation.

For an effective (directly usable) class we must choose an implementation. Let us use the array implementation illustrated at the beginning of the discussion of abstract data types:



The array will be called *representation*, and will have bounds 1 and *capacity*; the implementation also uses an integer, the attribute *count*, to mark the top of the stack.

Note that as we discover inheritance we will see how to write deferred classes that cover several possible implementations rather than just one. Even for a class that uses a particular implementation, for example by arrays as here, we will be able to inherit from the implementation class *ARRAY* rather than use it as a client (although some object-oriented developers will still prefer the client approach). For the moment, however, we can do without any inheritance-related technique.

Here is the class. Recall that if *a* is an array then the operation to assign value *x* to its *i*-th element is *a.put (x, i)*, and the value of its *i*-th element is given by *a.item (i)* or, equivalently, *a @ i*. If, as here, the bounds of the array are 1 and *capacity*, then *i* must in all cases lie between these bounds.

indexing

*description: "Stacks: Dispenser structures with a Last-In, First-Out %
%access policy, and a fixed maximum capacity"*

class STACK2 [G] creation

make

feature -- Initialization

make (n: INTEGER) is

– Allocate stack for a maximum of *n* elements

require

positive_capacity: n >= 0

do

capacity := n

!! representation.make (1, capacity)

ensure

capacity_set: capacity = n

array_allocated: representation /= Void

stack_empty: empty

end

feature -- Access

capacity: INTEGER

– Maximum number of stack elements

count: INTEGER

– Number of stack elements

item: G is

– Top element

require

not_empty: not empty – i.e. *count > 0*

do

Result := representation @ count

end

feature -- Status report

empty: BOOLEAN is

– Is stack empty?

do

Result := (count = 0)

ensure

empty_definition: Result = (count = 0)

end

On the export status
of *capacity* see ex-
cise E11.4, page 416.

full: BOOLEAN is

– Is stack full?

do

Result := (count = capacity)

ensure

full_definition: Result = (count = capacity)

end

feature -- Element change

put (x: G) is

– Add *x* on top

require

not_full: not full

– i.e. *count < capacity* in this representation

do

count := count + 1

representation.put (count, x)

ensure

not_empty: not empty

one_more_item: count = old count + 1

in_top_array_entry: representation @ count = x

end

remove is

– Remove top element

require

not_empty: not empty – i.e. *count > 0*

do

count := count - 1

ensure

not_full: not full

one_fewer: count = old count - 1

end

feature {NONE} -- Implementation

representation: ARRAY [G]

– The array used to hold the stack elements

invariant

... To be filled in later (see page 365) ...

end -- class STACK2

This class text illustrates the simplicity of working with assertions. It is complete except for the *invariant* clause, which will be added later in this chapter. Let us explore its various properties.

*Invariants are
introduced in
"CLASS INVARI-
ANTS", 11.8, page
364.*

This is the first full-fledged class of this chapter; not too far from what you will find in professional libraries of reusable object-oriented components such as the Base libraries. (Apart from the use of inheritance and a few extra features, what still distinguishes this class from its real-life counterparts is the absence of the *invariant* clause.)

On multiple feature clauses and exporting to *NONE* see “SELECTIVE EXPORTS AND INFORMATION HIDING”, 7.8, page 191.

Before studying the assertions, a general note about the structure of the class. As soon as a class has more than two or three features, it becomes essential to organize its features in a coherent way. The notation helps by providing the possibility of including multiple *feature* clauses. An earlier chapter introduced this facility as a way to specify a different export status for certain features, as done here for the last part of the class, labeled – *Implementation* to specify that feature *representation* is secret. But as already previewed in *STACK1* you can take advantage of multiple feature clauses even when the export status is the same. The purpose is to make the class easier to read, and easier to manage, by grouping features into general categories. After each *feature* keyword appears a comment (known as the Feature Clause Comment) defining the general role of the features that follow. The categories used in the example are those of *STACK1*, plus *Initialization* for the creation procedure.

“Feature clause header comments”, page 888.

The standard feature categories and associated Feature Clause Comments are part of the general rules for consistency and organization of reusable library classes. A more complete list appears in the chapter on style rules.

The imperative and the applicative

“Introducing an imperative view”, page 145.

The assertions of *STACK2* illustrate a fundamental concept of which we got a first glimpse when we studied the transition from abstract data types to classes: the difference between imperative and applicative views.

The assertions in *empty* and *full* may have caused you to raise an eyebrow. Here again is the text of *full*:

```
full: BOOLEAN is
    – Is stack full?

    do
        Result := (count = capacity)
    ensure
        full_definition: Result = (count = capacity)
    end
```

The postcondition expresses that *Result* has the same value as *count = capacity*. (Since both sides of the equality, the entity *Result* and the expression *count = capacity*, are boolean, this means that the function returns *true* if and only if *count* is equal to *capacity*.) But what is the point of writing this postcondition, one may ask, since the body of the routine (the *do* clause) says exactly the same thing through the instruction *Result := (count = capacity)*, whose only difference with the postcondition clause is its use of *=* rather than *=?* Is the postcondition not redundant?

Actually, there is a big difference between the two constructs, and no redundancy at all. The instruction *Result := (count = capacity)* is a command that we give to our virtual computer (the hardware-software machine) to change its state in a certain way; it performs an action. The assertion *Result = (count = capacity)* does not do anything; it specifies a property of the expected end state, as visible to the routine’s caller.

The instruction is *prescriptive*; the assertion is *descriptive*. The instruction describes the “how”; the assertion describes the “what”. The instruction is part of the implementation; the assertion is an element of specification.

The instruction is *imperative*; the assertion is *applicative*. These two terms emphasize the fundamental difference between the worlds of computing and mathematics:

- Computer operations may change the state of the hardware-software machine. Instructions of common programming languages are commands (imperative constructs) directing the machine to execute such operations.
- Mathematical reasoning never changes anything; as noted in the presentation of abstract data types, taking the square root of the number 2 does not change that number. Mathematics instead describes how to use properties of known objects, such as the number 2, to infer properties of others, such as $\sqrt{2}$, obtained from the former by *applying* (hence the name) certain mathematical derivations such as square root.

That the two notations are so close in our example — assignment *:=* and equality *=* — should not obscure this fundamental difference. The assertion describes an intended result, and the instruction (the loop body) prescribes a particular way to achieve that result. Someone using the class to write a client module will typically be interested in the assertion but not in the implementation.

The reason for the closeness of notations for assignment and equality is that assignment is indeed in many cases the straightforward way to achieve equality; in our example the chosen implementation, *Result := (count = capacity)*, is indeed the obvious one. But as soon as we move on to more advanced examples the conceptual difference between the specification and the implementation will be much larger; even in the simple case of a function to compute the square root of a real number *x*, where the postcondition is just something like *abs (Result ^ 2 - x) <= tolerance* with *abs* denoting absolute value and *tolerance* a tolerance value, the instructions in the function’s body will be far less trivial since they have to implement a general algorithm for the computation of square roots.

Even for *put* in class *STACK2*, the same specification could have led to different implementations, although the differences are minor; for example the body could be

```
if count = capacity then Result := True else Result := False end
perhaps simplified (thanks to the rules of default initialization) into
if count = capacity then Result := True end
```

So the presence of related elements in the body and the postcondition is not evidence of redundancy; it is evidence of consistency between the implementation and the specification — that is to say, of correctness as defined at the beginning of this chapter.

In passing, we have encountered a property of assertions that will merit further development: their relevance for authors of client classes, whom we should not ask to read routine implementations, but who need a more abstract description of the routine's role. This idea will lead to the notion of **short form** discussed later in this chapter as the basic class documentation mechanism.

See "Including functions in assertions", page 461.

A caveat: for practical reasons we will allow assertions to include some seemingly imperative elements (functions). This issue will be explored at the end of this chapter.

As a summary of this discussion it is useful to list the words that have been used to contrast the two categories of software elements:

Implementation	Specification
Instruction	Expression
How	What
Imperative	Applicative
Prescription	Description

The imperative-applicative opposition

A note on empty structures

The precondition of the creation procedure *make* in class *STACK1* requires a comment. It states $n \geq 0$, hence allowing empty stacks. If n is zero, *make* will call the creation procedure for arrays, also named *make*, with arguments *1* and *0* for the lower and upper bounds respectively. This is not an error, but follows from a convention regarding *ARRAY*'s creation procedure: using a first argument greater than the second by one creates an empty array.

A zero n for a stack, or a first creation argument greater than the second for an array, is not wrong but simply means that this particular stack or array should be empty. An error would only occur out of a call attempting to access an element from the structure, for example a *put* for the stack or an *item* for the array, both of whose preconditions will always be false for an empty structure ("my customer is always wrong").

When you define a general data structure such as a stack or array, you should determine whether the case of an empty structure is conceptually meaningful. In some cases it is not: for example most definitions of the notion of *tree* start from the assumption that there is at least one node, the root. But if the empty case raises no logical impossibility, as with arrays and stacks, you should plan for it in the design of your data structure, acknowledging that clients will, every once in a while, create empty instances, and should not suffer for it. An application system may for example need a stack for n elements, where n is an upper bound on the number of elements to be stacked, computed by the application just before it creates the stack; in some runs that number may be zero. This is not an error, simply an extreme case.

The array mechanism of Algol W provides a counter-example. When a dynamically allocated array has an empty range, the program terminates in error — even if it was a perfectly valid array which simply happened to be empty on that particular run. This is too restrictive: an array with zero size is valid, it simply does not allow access to any element.

Precondition design: tolerant or demanding?

Central to Design by Contract is the idea, expressed as the Non-Redundancy principle, that for any consistency condition that could jeopardize a routine's proper functioning you should assign enforcement of this condition to only one of the two partners in the contract.

The principle was on page 344.

Which one? In each case you have two possibilities:

- Either you assign the responsibility to clients, in which case the condition will appear as part of the routine's precondition.
- Or you appoint the supplier, in which case the condition will appear in a conditional instruction of the form *if condition then ...*, or an equivalent control structure, in the routine's body.

We can call the first attitude *demanding* and the second one *tolerant*. The *STACK2* class illustrates the demanding style: a tolerant version of the class would have routines with no preconditions, such as

```
remove is
do
    -- Remove top element
    if empty then
        print ("Error: attempt to pop an empty stack")
    else
        count := count - 1
    end
end
```

Warning: not the recommended style.

In the analogy with human contracts we can think of the demanding style as characterizing an experienced contractor who expects his clients to "do their homework" before calling on him; he has no trouble finding business, and will reject requests that appear too broad or unreasonable. The tolerant style evokes the image of a freshly established consulting practice, whose owner is so desperate for business that he will take anything, having put in his driveway a big sign:



Which is the better style? To a certain extent this is a matter of personal choice (as opposed to the Non-Redundancy principle, which was absolute in stating that it is *never* acceptable to deal with a correctness condition on both the client and supplier sides). A strong case can be made, however, for the demanding style illustrated by *STACK2*, especially in the case of software meant to be reusable — and in O-O development we should always write our software with the goal of ultimately making it reusable.

At first the tolerant style might appear better for both reusability and reliability; after all the demanding approach appears to put more responsibility on the clients, and there are typically many clients for a single supplier — even more so for a reusable class. Is it not preferable then, to let the supplier take care of the correctness conditions once and for all, rather than require every client to do it for itself?

If we look more closely at the issue this reasoning does not hold. The correctness conditions describe what the routine requires to be able to do its job properly. The tolerant *remove* on the facing page is a good counter-example: what can a poor stack-popping routine do for an empty stack? It makes a brave attempt by outputting an error message, but this is clearly inadequate: a specialized utility module such as a stack handler has no business messing up the system's user output. We could try something more sophisticated, but *remove* simply does not have the proper context: the focus of class *STACK2* is too narrow to determine what to do in the case of an empty stack. **Only the client** — a module using stacks in some application, for example the parsing module in a compiler — has enough information to decide what an attempt to pop an empty stack really means: is it a normal although useless request that we should simply ignore, executing a null operation? Or is it an error, and if so, how should we handle it: raise an exception, correct the situation before trying again, or (the least likely answer) output a user-visible error message?

In the square root example, you may remember the fictitious routine text quoted in the discussion preceding the Non-Redundancy principle:

```
if x < 0 then
    "Handle the error, somehow"
else
    "Proceed with normal square root computation"
end
```

Page 344.

The operative word is “*somehow*”. The **then** clause is incarnation more than software: there is really no good general-purpose technique for handling the $x < 0$ case. Here again a general-purpose routine has no clue. Only the client author can know what the call means in this case — an error in the software, a case in which the expected result is 0, a reason to trigger an exception...

In this case as in the attempt at a tolerant *remove*, the position of the routine is not unlike that of a postman asked to deliver a postcard with no delivery address and no return address: the case falls outside of the contract, and there is no good way to decide what to do.

In the spirit of Design by Contract, the demanding approach to precondition design does not attempt to produce routines that are all things to all clients. Instead, it insists that each routine do a well-defined job and do it well (correctly, efficiently, generally enough to be reusable by many clients...) and specify clearly what cases it cannot handle. In fact you cannot hope that the routine will do its job well *unless* you have carefully circumscribed that job. A factum routine, which wants to do a computation and check for abnormal cases and take corrective actions and notify the client and produce a result anyway, will most likely fail to fulfill any of these goals properly.

The routine author does not try to outsmart his clients; if he is not sure of what the routine is supposed to do in a certain abnormal situation, he excludes it explicitly through the precondition. This attitude is more generally a consequence of the overall theme in this book: building software systems as sets of modules that mind their own business.

If you read the supplementary mathematical section in the chapter on abstract data types, you may have noted the similarity between the present discussion and the arguments for using partial functions in the mathematical model, rather than special error values such as *undefined*. The two ideas are indeed very close, and Design by Contract is in part the application to software construction of the concept of partial function, so remarkably flexible and powerful in formal specification.

A word of caution: the demanding approach is only applicable if the preconditions remain reasonable. Otherwise the job of writing a module would become easy: start every routine with **require False** so that, as we have seen, any routine body will be correct. What does “reasonable” concretely mean for the precondition of a routine? Here is a more precise characterization:

Reasonable Precondition principle

Every routine precondition (in a “demanding” design approach) must satisfy the following requirements:

- The precondition appears in the official documentation distributed to authors of client modules.
- It is possible to justify the need for the precondition in terms of the specification only.

^{“Alternatives to partial functions”, page 151.}

The first requirement will be supported by the notion of short form studied later in this chapter. The second requirement excludes restrictions meant only for the supplier's convenience in implementing the routine. For example when you want to pop a stack the precondition *not empty* is a logical requirement that can be justified "in terms of the specification only", through the simple observation that in an empty stack there is nothing to pop; and when you want to compute the real square root of a number, the precondition $x \geq 0$ is a direct result of the mathematical property that negative real numbers do not have real square roots.

Some restrictions may arise from the general kind of implementation selected. For example the presence of *require not full* as precondition to the push operation *put* in *STACK2* is due to the decision of using an array for the implementation of stacks. But such a case does not violate the principle, as the bounded nature of *STACK2* stacks has been made part of the specification: the class does not claim to represent arbitrary stacks, but only stacks of finite maximum capacity (as expressed for example in the *indexing* clause of the class). The abstract data type serving as specification of this class is not the most general notion of stack, but the notion of bounded stack.

In general, it is desirable to avoid bounded structures; even a stack implemented by arrays can use array resizing. This is the case with the most commonly used stack class in the Base libraries, which follows the *STACK2* style but without a notion of *capacity*; a stack that overflows its current capacity resizes itself silently to accommodate the new elements.

Preconditions and export status

You may have noted the need for a supplementary requirement on preconditions, which does not figure in the Reasonable Precondition principle: to be satisfiable by the clients, the precondition must not use features that are hidden from the clients as a result of export restrictions.

Assume for example the following situation:

```
-- Warning: this is an invalid class, for purposes of illustration only.
class SNEAKY Feature
  tricky is
  require
    accredited
  do
    ...
  end
feature {NONE}
  accredited: BOOLEAN is do ... end
end -- class SNEAKY
```

The specification for *tricky* states that any call to that procedure must satisfy the condition expressed by the boolean function *accredited*. But whereas the class exports *tricky* to all clients, it keeps *accredited* secret, so that clients have no way of finding out,

before a call, whether the call is indeed correct. This clearly unacceptable situation is akin, in human contracts, to a deal in which the supplier would impose some conditions not stated explicitly in the contract, and hence could reject a client's request as incorrect without giving the client any way to determine in advance whether it is correct.

The reason why the Reasonable Precondition principle does not cover such cases is that here a methodological principle does not suffice: we need a language rule to be enforced by compilers, not left to the decision of developers.

The rule must take into account all possible export situations, not just those illustrated above in which a feature is available to all clients (*tricky*) or to no client (*accredited*). As you will recall from the discussion of information hiding, it is also possible to make a feature available to some clients only, by declaring it in a feature clause appearing as *feature {A, B, ...}*, which makes it available only to *A, B, ...* and their descendants. Hence the language rule:

Precondition Availability rule

Every feature appearing in the precondition of a routine must be available to every client to which the routine is available.

With this rule every client that is in a position to call the feature will also be in a position to check for its precondition. The rule makes class *SNEAKY* invalid, since *tricky* is generally exported (available to all clients); you can turn it into a valid class by making *accredited* also generally exported. If *tricky* had appeared in a feature clause starting with *feature {A, B, C}*, then *accredited* would have to be exported at least to *A, B* and *C* (by appearing in the same feature clause as *tricky*, or by appearing in a clause of the form *feature {A, B, C}, or feature {A, B, C, D, ...}*, or just *feature*). Any violation of this rule is a compile-time error. Class *SNEAKY*, for example, will be rejected by the compiler.

There is no such rule for postconditions. It is not an error for some clauses of a postcondition clause to refer to secret features, or features that are not as broadly exported as the enclosing routine; this simply means that you are expressing properties of the routine's effect that are not directly usable by clients. This was the case with the *put* procedure in *STACK2*, which had the form

```
put (x: G) is
  -- Add x on top
  require
    not full
  do
    ...
  ensure
    ... Other clauses ...
  in_top_entry: representation @ count = x
end
```

For the Other clauses see page 351.

SELECTIVE EXPORTS AND INFORMATION HIDING, 78, page 191.

The last postcondition clause indicates that the array entry at index *count* contains the element just pushed. This is an implementation property; even though *put* is generally available (exported to all clients), array *representation* is secret. But there is nothing wrong with the postcondition; it simply includes, along with properties that are directly useful to clients (the “*Other Clauses*”), one that is only meaningful for someone who reads the entire class text. Such secret clauses will not appear in the “short” form of the class — the documentation for client authors.

A tolerant module

(On first reading you may skip this section or just look through it quickly.)

For filters of the first kind see “Assertions are not on input checking mechanism”, page 346.

The simple but unprotected basic modules may not be robust enough for use by arbitrary clients. In some cases there will be a need for new classes to serve as filters, interposed not between the software and the external world (as with filters of the kind discussed earlier in this chapter) but between software and other software: possibly careless clients on one side, unprotected classes on the other.

Although we have seen that this is generally not the right approach, it is useful to examine how classes will look if we do decide to use the tolerant style in a specific case. Class *STACK3*, appearing next, illustrates the idea. Because the class needs to set integer error codes, it is convenient to rely on a property of the notation that has not been introduced yet: “unique” integer constants. If you declare a set of attributes as

a, b, c, ...: INTEGER is unique

“UNIQUE VALUES”, 18.6, page 654.

the effect is to define *a, b, c, ...* as integer constants with consecutive positive values. These values will be assigned by the compiler, and are guaranteed to be different for all constants thus declared, relieving you of having to invent separate codes. By convention, constant attributes such as these have names beginning with an upper-case letter, with the rest in lower case, as in *Underflow*.

Here, using this technique, is a tolerant version of our earlier stack class. Make sure to note that this class text (which you may just skim through on first reading) is included here only to make sure you understand the tolerant style; it is **not** an example of the generally recommended design — for reasons that will be discussed below, but will probably be clear enough as you browse through the text.

Indexing

*description: “Stacks: Dispenser structures with a Last-In, First-Out %
%access policy, and a fixed maximum capacity; %
%tolerant version, setting an error code in case %
%of impossible operations.”*

class STACK3 [G] creation
make

```
feature -- Initialization
  make (n: INTEGER) is
    -- Allocate stack for a maximum of n elements if n > 0;
    -- otherwise set error to Negative_size.
    -- No precondition!
  do
    if capacity >= 0 then
      capacity := n
      !! representation .make (capacity)
    else
      error := Negative_size
    end
  ensure
    error_code_if_impossible: (n < 0) = (error = Negative_size)
    no_error_if_possible: (n >= 0) = (error = 0)
    capacity_set_if_no_error: (error = 0) implies (capacity = n)
    allocated_if_no_error: (error = 0) implies (representation /= Void)
  end
end

feature -- Access
  item: G is
    -- Top element if present; otherwise the type's default value.
    -- with error set to Underflow.
    -- No precondition!
  do
    if not empty then
      check representation /= Void end
      Result := representation.item
      error := 0
    else
      error := Underflow
      -- In this case the result is the default value
    end
  ensure
    error_code_if_impossible: (old empty) = (error = Underflow)
    no_error_if_possible: (not (old empty)) = (error = 0)
  end
end

feature -- Status report
  empty: BOOLEAN is
    -- Number of stack elements
  do
    Result := (capacity = 0) or else representation.empty
  end
end
```

```

error: INTEGER
-- Error indicator, set by various features to a non-zero value
-- If they cannot do their job
full: BOOLEAN is
-- Number of stack elements
do
    Result := (capacity = 0) or else representation.full
end
Overflow, Underflow, Negative_size: INTEGER is unique
-- Possible error codes
feature -- Element change
    put (x: G) is
        -- Add x on top if possible; otherwise set error code.
        -- No precondition!
    do
        if full then
            error := Overflow
        else
            check representation /= Void end
            representation.put (x); error := 0
        end
    ensure
        error_code_if_impossible: (old full) = (error = Overflow)
        no_error_if_possible: (not old full) = (error = 0)
        not_empty_if_no_error: (error = 0) implies not empty
        added_to_top_if_no_error: (error = 0) implies item = x
        one_more_item_if_no_error: (error = 0) implies count = old count + 1
    end
remove is
    -- Remove top element if possible; otherwise set error.
    -- No precondition!
    do
        if empty then
            error := Underflow
        else
            check representation /= Void end
            representation.remove
            error := 0
        end
    ensure
        error_code_if_impossible: (old empty) = (error = Underflow)
        no_error_if_possible: (not old empty) = (error = 0)
        not_full_if_no_error: (error = 0) implies not full
        one_fewer_item_if_no_error: (error = 0) implies count = old count - 1
    end

```

```

feature {NONE} -- Implementation
    representation: STACK2 [G]
        -- The unprotected stack used as implementation
    capacity: INTEGER
        -- The maximum number of stack elements
end -- class STACK3

```

The operations of this class have no preconditions (or, more accurately, have *True* as their preconditions). For those that may result in abnormal situations, the postcondition has been refined to distinguish between correct and erroneous processing. An operation such as *s.remove*, where *s* is a *STACK3*, will set *s.error* to 0 or to *Underflow* (which, from the rules on unique values, is known to be strictly positive) and, in the latter case, do nothing else. It is still the caller's responsibility to check for *s.error* after the call. As noted, a general-purpose module such as *STACK3* has no way to decide what to do in the case of an erroneous popping attempt: produce an error message, take corrective action...

Such filter modules serve to separate algorithmic techniques to deal with normal cases and techniques for handling errors. This is the distinction between correctness and robustness explained at the beginning of this book: writing a module that performs correctly in legal cases is one task; making sure that other cases are also processed decently is another. Both are necessary, but they should be handled separately. Failure to do so is one of the principal reasons why so many software systems are hopelessly complex: any algorithm that does anything useful also takes care of checking that it is applicable, and for good measure tries to handle the cases in which it is not. Such software soon mushrooms into a total mess.

A few technical comments apply to this example:

- An instance of *STACK3* is not an array but a structure containing a reference (*representation*) to an instance of *STACK2*. Itself containing a reference to an array. These two indirections, detrimental to efficiency, can be avoided through inheritance as studied in later chapters.
- The boolean operator *or else* is similar to *or* but ignores the second operand if it does not affect the result and trying to evaluate it could cause trouble.
- The *check* instruction used in *put* and *remove* serves to state that a certain assertion is satisfied. It will be studied later in this chapter.

Finally, you will have noted the heaviness of *STACK3*, especially if you compare it to the simplicity that *STACK2* achieves with its precondition. *STACK3* is good evidence that a tolerant style may lead to uselessly complex software. The demanding style, in contrast, follows from the general spirit of Design by Contract. Trying to handle all possible (and impossible) cases is not necessarily the best way to help your clients. If instead you build classes that impose possibly strict but reasonable usage conditions, and describe these conditions precisely as part of the official documentation for the class, you actually make life easier for the clients. This has been called the *tough love* approach: you can often serve your clients better by being more restrictive.

See "A REVIEW OF EXTERNAL FACTORS", 1.2, page 4.

Or, or else see "Non-strict boolean operators", page 454.

Better an efficient supplier that states its functionally justified limitations than a overzealous one that tries to second-guess its clients, making possibly inappropriate decisions for abnormal cases, and sacrificing simplicity and efficiency.

For modules whose clients are other software modules, the demanding approach is usually the right one. A possible exception is the case of modules intended for clients whose authors use a non-O-O language and may not have understood the basic concepts of Design by Contract.

"Assertions are not an input checking mechanism" - page 346.

The tolerant approach remains useful for software elements that deal not with other software elements but with data coming from the outside world, such as user input, or sensor data. Then, as noted earlier, filter modules are often necessary to separate the actual processing modules (the physicists in our metaphor) from those which simply quality data and reject anything that is not appropriate (the guards). This separation of concerns is essential for maintaining the simplicity of software elements on both sides. *STACKS* provides an idea of what such modules may look like.

11.8 CLASS INVARIANTS

Preconditions and postconditions describe the properties of individual routines. There is also a need for expressing global properties of the instances of a class, which must be preserved by all routines. Such properties will make up the class invariant, capturing the deeper semantic properties and integrity constraints characterizing a class.

Definition and example

Consider again the earlier implementation of stacks by arrays, the one without the protections (*STACK2*):

```
class STACK2 [G] creation
  make
  feature
    ... make, empty, full, item, put, remove ...
  capacity: INTEGER
  count: INTEGER
  feature (NONE) -- Implementation
  representation: ARRAY [G]
end
```

The attributes of the class — array *representation* and integers *capacity* and *count* — constitute the stack representation. Although routine preconditions and postconditions, given earlier, express some of the semantic properties of stacks, they fail to express other important consistency properties linking the attributes. For example, *count* should always remain between 0 and *capacity*:

0 <= count; count <= capacity
(implying also that *capacity >= 0*), and *capacity* should be the array size:

capacity = representation.capacity

A class invariant is such an assertion, expressing general consistency constraints that apply to every class instance as a whole; this is different from preconditions and postconditions, which characterize individual routines.

The above assertions involve only attributes. Invariants may also express the semantic relations between functions, or between functions and attributes. For example the invariant for *STACK2* may include the following property describing the connection between *empty* and *count*:

empty = (count = 0)

In this example, the invariant assertion links an attribute and a function; it is not particularly interesting as it merely repeats an assertion that appears in the postcondition of the function (there *empty*). More useful assertions are those which involve either only attributes, as above, or more than one function.

Here is another typical example. Assume — in line with previous examples dealing with the notion of bank account — that we have a class *BANK_ACCOUNT* with features *deposits_list*, *withdrawals_list* and *balance*. Then the invariant for such a class could include a clause of the form:

consistent_balance: deposits_list.total - withdrawals_list.total = balance

where the function *total* gives the cumulated value of a list of operations (deposits or withdrawals). This states the basic consistency condition between the values accessible through features *deposits_list*, *withdrawals_list* and *balance*.

Form and properties of class invariants

Syntactically, a class invariant is an assertion, appearing in the *invariant* clause of the class, after the features and just before the *end*, as in

```
class STACK4 [G] creation
  ... As in STACK2 ...
  feature
    ... As in STACK2 ...
  invariant
    count_non_negative: 0 <= count
    count_bounded: count <= capacity
    consistent_with_array_size: capacity = representation.capacity
    empty_if_no_elements: empty = (count = 0)
    item_at_top: (count > 0) implies (representation.item (count) = item)
  end
```

An invariant for a class *C* is a set of assertions that every instance of *C* will satisfy at all "stable" times. Stable times are those in which the instance is in an observable state:


This example was first discussed in "UniformAccess", page 35. It will serve again to illustrate persistence issues: "Correction", page 1045.

For the features of STACK2 see page 350.

- On instance creation, that is to say after execution of $!!\ a$ or $!!\ a.mkde(\dots)$, where a is of type C .

- Before and after every remote call $a.r(\dots)$ to a routine r of the class.

The following figure, showing the life of an object, helps put the notions of invariant and stable time in place.



Life as an object, to tell the truth, is not that thrilling (in case you ever wondered). At the beginning — left of the figure — you do not exist. You are begot by a creation instruction $!!\ a$ or $!!\ a.mkde(\dots)$, or a *clone*, and reach your first station in life. Then things get quite boring: through some reference a , clients use you, one after the other, by applying operations of the form $a.f(\dots)$ where f is a feature of your generating class. And so on forever, or at least until execution terminates.

The invariant is the characteristic property of the states represented by gray squares in the figure — S1 etc. These are the “stable times” mentioned above: those at which the object is observable from the outside, in the sense that a client can apply a feature to it. They include:

- The state that results from the creation of an object (S1 in the figure).
 - The states immediately before and after a call of the form $a.some_routine(\dots)$ executed by a client.
- Here the context is sequential computation, but the ideas will transpose to concurrent systems in a later chapter.

Chapter 30 covers concurrency; see in particular “Concurrent accesses to an object”, page 30c.

An invariant that varies

In spite of its name, the invariant does not need to be satisfied at all times, although in the *STACK4* example it does remain true after the initial creation. In the more general case, it is perfectly acceptable for a procedure g to begin by trying to work towards its goal — its postcondition — and in the process to destroy the invariant (as in human affairs, trying to do something useful may disrupt the established order of things), then it spends the second part of its execution scurrying to restore the invariant without losing too much of whatever ground has been gained. At some intermediate stages, such as the instant marked ■ in the figure, the invariant will not hold; this is fine as long as the procedure reestablishes the invariant before terminating its execution.

Who must preserve the invariant?

Qualified calls, of the form $a.f(\dots)$, executed on behalf of a client, are the only ones that must always start from a state satisfying the invariant and leave a state satisfying the invariant: there is no such rule for unqualified calls of the form $f(\dots)$, which are not directly executed by clients but only serve as auxiliary tools for carrying out the needs of qualified calls. As a consequence, the obligation to maintain the invariant applies only to the body of features that are exported either generally or selectively: a secret feature — one that is available to no client — is not affected by the invariant.

From this discussion follows the rule that precisely defines when an assertion is a correct invariant for a class:

Invariant rule
An assertion I is a correct class invariant for a class C if and only if it meets the following two conditions:
E1 • Every creation procedure of C , when applied to arguments satisfying its precondition in a state where the attributes have their default values, yields a state satisfying I .
E2 • Every exported routine of the class, when applied to arguments and a state satisfying both I and the routine’s precondition, yields a state satisfying I .

Note that in this rule:

- Every class is considered to have a creation procedure, defined as a null operation if not explicitly specified.
- The state of an object is defined by all its fields (the values of the class attributes for this particular instance).
- The precondition of a routine may involve the initial state and the arguments.

- The postcondition may only involve the final state, the initial state (through the **old** notation) and, in the case of a function, the returned value, given by the predefined entity *Result*.
- The invariant may only involve the state.

Assertions may use functions, but such functions are an indirect way of referring to the attributes — to the state.

A mathematical statement of the Invariant rule appears later in this chapter.

This is the topic of exercise §11.8, page 410. The reasoning for preconditions and postconditions was in “Assertions are not control structures”, page 347.

You can use the Invariant rule as a basis for answering a question that comes up in light of earlier discussions: what would it mean if an invariant clause turned out to be violated during system execution? We saw before that a precondition violation signals an error (a “bug”) in the client, a postcondition violation an error in the supplier. The answer will be for invariants as for postconditions; you have all the elements for deriving this property by yourself.

The role of class invariants in software engineering

Property E2 indicates that we may consider the invariant as being implicitly added (and) to both the precondition and postcondition of every exported routine. So in principle the notion of invariant is superfluous: we could do without it by enriching the preconditions and postconditions of all routines in the class.

Such a transformation is of course not desirable. It would complicate the routine texts; but more importantly, we would lose the deeper meaning of the invariant, which transcends individual routines and applies to the class as a whole. One should in fact consider that the invariant applies not only to the routines actually written in the class, but also to any ones that might be added later, thus serving as control over future evolution of the class. This will be reflected in the inheritance rules.

In the view of software development introduced at the beginning of this book, we accept that change is inevitable, and we try to control it. Some aspects of a software system, and of its individual components — classes — may be expected to change faster than others. Adding, removing or changing features, in particular, is a frequent and normal event. In this volatile process one will want to cling to properties that, although they may change too — for we can hardly guarantee that any aspect of a system will remain set for eternity — will change far less often. Invariants, because they capture the fundamental semantic constraints applying to a class, play this role.

The *STACK2* example illustrates the basic ideas, but to appreciate the full power of the concept of invariant you should be on the lookout for further examples of invariants in the rest of this book. To me the notion of the invariant is one of the most illuminating concepts that can be learned from the object-oriented method. Only when I have derived the invariant (for a class that I write) or read and understood it (for someone else’s class) do I feel that I know what the class is about.

Invariants and contracting

Invariants have a clear interpretation in the contract metaphor. Human contracts often contain references to general clauses or regulations that apply to all contracts within a certain category; think of a city’s zoning regulations, which apply to all house-building contracts. Invariants play a similar role for software contracts: the invariant of a class affects all the contracts between a routine of the class and a client.

Let us probe further. It was noted above that we may consider the invariant as being added to both the precondition and postcondition of every exported routine. Let *body* be the body of a routine (the set of instructions in its **do** clause), *pre* its precondition, *post* its postcondition and *INV* the class invariant. The correctness requirement on the routine may be expressed, using the notation introduced earlier in this chapter, as:

{*INV* and *pre*} *body* {*INV* and *post*}

(As you will remember this means: any execution of *body*, started in any state in which *INV* and *pre* both hold, will terminate in a state in which both *INV* and *post* hold.)

For the supplier author — the person who writes *body* — is the invariant good news or bad news, that is to say, does it make the job easier or harder?

The answer, as you will have figured out from the earlier discussion, is: both. Remember our lazy job applicant, who wanted a *strong* precondition and a *weak* postcondition. Here adding *INV* makes stronger or equal both the precondition and the postcondition. (From the rules of logic, *a* and *b* always implies *a*, that is to say, is stronger than or equal to *a*.) So, if you are in charge of implementing the *body*, the invariant:

- Makes your job easier: in addition to the official precondition *pre*, you may assume that the initial state satisfies *INV*, further restricting the set of cases that you must handle.
- Makes your job harder: in addition to your official postcondition *pos*, you must ensure that the final state satisfies *INV*.

These observations are consistent with the view of the invariant as a general consistency condition that applies to the class as a whole, and hence to all of its routines. As the author of such a routine, you have the benefit of being permitted to take this condition for granted at the start of the routine; but you have the obligation to ensure that the routine will satisfy it again on termination — so that the next routine to be executed on the same object can in turn take it for granted.

The class *BANK_ACCOUNT* mentioned above, with the invariant clause

deposits_list.total – withdrawals_list.total = balance

provides a good example. If you have to add a routine to the class, this clause gives you the guarantee that the features *deposits_list*, *withdrawals_list* and *balance* have consistent values, so you do not need to check this property (and then, as we have seen, you **must not** check it). But it also means that you must write the routine so that, whatever else it does, it will leave the object in a state that again satisfies the property. So a procedure *withdraw*,

The notation was defined on page 355.

used to record a withdrawal operation, should not just update *withdrawals.list*; it must also, if *balance* is an attribute, update the value of *balance* to take the withdrawal into account and restore the invariant, enabling any other routine called later on the same object to benefit from the same original assumption that facilitated the work of *withdraw*.

Rather than an attribute, *balance* could be a function, whose body computes and returns the value of *deposits.list.total* – *withdrawals.list.total*; in this case procedure *withdraw* does not need to do anything special to maintain the invariant. The ability to switch at will between the two representations without affecting the client is an illustration of the principle of Uniform Access.

See “Uniform Access”, page 55.

This example shows the idea of class invariant as a transposition to software of one of the rules of polite behavior: that if you use a shared facility — say an office kitchen — you should leave it for others, after each use, in the state in which you would like to find it when you start.

11.9 WHEN IS A CLASS CORRECT?

If you prefer to skip the theory, you should read the next section, “CONSTRUCTING A CLASS”, page 375.

Although we still have to see a few more constructs involving assertions, it is useful to take a brief pause and examine some of the implications of what we have learned about preconditions, postconditions and invariants. This section does not introduce any new constructs, but describes some of the theoretical background. Even on your first reading I think you should get familiar with these ideas as they are central to a proper understanding of the method, and will be precious when we try to figure out how to use inheritance well.

The correctness of a class

With preconditions, postconditions and invariants, we can now define precisely what it means for a class to be correct.

The basis for the answer appeared at the beginning of this chapter: a class, like any other software element, is correct or incorrect not by itself but with respect to a specification. By introducing preconditions, postconditions and invariants we have given ourselves a way to include some of the specification in the class text itself. This provides a basis against which to assess correctness: the class is correct if and only if its implementation, as given by the routine bodies, is consistent with the preconditions, postconditions and invariant.

The notation $(P) A (Q)$ introduced at the beginning of this chapter helps express this precisely. Remember that the meaning of such a correctness formula is: whenever A is executed in a state satisfying P , the execution will terminate in a state satisfying Q .

Let C be a class. *INV* its class invariant. For any routine r of the class, call $pre_r(x_i)$ and $post_r(x_i)$ its precondition and postcondition; x_i denotes the possible arguments of r , to which both the precondition and the postcondition may refer. (If the precondition or postcondition is missing from the routine text, then pre_r or $post_r$ is just *True*.) Call *Body_r* the body of routine r .

Finally, let *Default_C* be the assertion expressing that the attributes of C have the default values of their types. For example *Default_{STACK}* referring to the earlier stack class, is the assertion

representation = Void
capacity = 0
count = 0


These notations permit a general definition of class correctness:

Definition: class correctness

A class is correct with respect to its assertions if and only if:

- C1 • For any valid set of arguments x_p to a creation procedure p :
 $(Default_C \text{ and } pre_p(x_p)) \text{ Body}_p\{post_p(x_p) \text{ and } INV\}$
- C2 • For every exported routine r and any set of valid arguments x_i :
 $(pre_r(x_i) \text{ and } INV) \text{ Body}_r\{post_r(x_i) \text{ and } INV\}$

This rule — previewed informally in the *BANK_ACCOUNT* example — is a mathematical statement of the earlier informal diagram showing the lifecycle of a typical object, which is worth looking at again:



The life of an object

(This figure first appeared on page 366)

Condition C1 means that any creation procedure (such as *make* in the figure), when called with its precondition satisfied, must yield an initial state (S1 in the figure) that satisfies the invariant and the procedure’s postcondition. Condition C2 expresses that any exported routine r (such as f or g in the figure), if called in a state (S1, S2 or S3) satisfying both its precondition and the invariant, must terminate in a state that satisfies both its postcondition and the invariant.

If we focus on invariants, we may look at the preceding definition of class correctness as working by induction on the set of instances of a class. Rule *C1* is the base step of the induction, stating that the invariant holds for all newborn objects — those which directly result from a creation instruction. Rule *C2* is the induction step, through which we determine that if a certain generation of instances satisfies the invariant, then the next generation — the set of instances obtained by applying exported features to the members of the current generation — will also satisfy it. Since by starting from newborn objects and going from generation to generation through exported features we obtain all possible instances of the class, the mechanism enables us to determine that all instances satisfy the invariant.

Two practical observations:

- If the class has no **creation** clause, we may consider that it has a single implicit creation procedure *nothing* with an empty body. Applying rule *C1* to *Nothing* then means that *Default_C* must imply *INV*: the default values must satisfy the invariant.
- A requirement of the form $\{P\} A \{Q\}$ does not commit *A* in any way for cases in which *P* is not initially satisfied. So the notation is in line with the property discussed in detail earlier in this chapter: the contract is not binding on the routine if the client fails to observe its part of the deal. Accordingly, the definition of class correctness leaves the routines of the class free to do as they please for any call that violates the precondition or the invariant.

What has just been described is how to *define* the correctness of a class. In practice, we may also want to *check* whether a given class is indeed correct. This issue will be discussed later in this chapter.

The role of creation procedures

The discussion of invariants yields a better understanding of the notion of creation procedure.

A class invariant expresses the set of properties that objects (instances of the class) must satisfy in what has been called the stable moments of their lifetime. In particular, these properties must hold upon instance creation.

The standard object allocation mechanism initializes fields to the default values of the corresponding attribute types; these values may or may not satisfy the invariant. If not, a specific creation procedure is required: it should set the values of the attributes so as to satisfy the invariant. So creation may be seen as the operation that ensures that all instances of a class start their lives in a correct mode — one in which the invariant is satisfied.

The first presentation of creation procedures introduced them as a way to answer a more mundane (and obvious) question: how do I override the default initialization rules if they do not suit me for a particular class, or if I want to provide my clients with more than one initialization mechanism? But with the introduction of invariants and the theoretical discussion summarized by rule *C1*, we also see the more profound role of creation procedures: they are here to make sure that any instance of the class, when it starts its life, already satisfies the fundamental rules of its caste — the class invariant.

See “CREATION PROCEDURES”, 8.4, page 238, in particular “Rules on creation procedures”, page 238

Arrays revisited

The library class *ARRAY* was sketched in the previous chapter. Only now, however, are we in a position to give its definition properly. The notion of array fundamentally requires preconditions, postconditions and an invariant.

See “ARRAYS”, 10.4, page 325.

Here is a better sketch with assertions. Preconditions express the basic requirement on array access and modification: indices should be in the permitted range. The invariant shows the relation between *count*, *lower* and *upper*; it would allow *count* to be implemented as a function rather than an attribute.

indexing

description: “Sequences of values, all of the same type or of a conforming one, % %accessible through integer indices in a contiguous interval”

class ARRAY [G] **creation**

make

feature -- Initialization

make (*minindex*, *maxindex*: INTEGER) **is**

– Allocate array with bounds *minindex* and *maxindex*

– (empty if *minindex* > *maxindex*).

require

meaningful_bounds: *maxindex* >= *minindex* – 1

do

...

ensure

exact_bounds_if_non_empty: (*maxindex* >= *minindex*) **implies**

((*lower* = *minindex*) **and** (*upper* = *maxindex*))

conventions_if_empty: (*maxindex* < *minindex*) **implies**

((*lower* = 1) **and** (*upper* = 0))

end

feature -- Access

lower, *upper*, *count*: INTEGER

– Minimum and maximum legal indices; array size.

infix “@”, *item* (i: INTEGER): G **is**

– Entry of index *i*

require

index_not_too_small: *lower* <= *i*

index_not_too_large: *i* <= *upper*

do ... **end**

feature — Element change

```
put (v: G; i: INTEGER) is
  — Assign v to the entry of index i
  require
    index_not_too_small: lower <= i
    index_not_too_large: i < upper
  do
    ...
  ensure
    element_replaced: item(i) = v
  end
invariant
  consistent_count: count = upper - lower + 1
  non_negative_count: count >= 0
end -- class ARRAY
```

The only part left blank is the implementation of routines *item* and *put*. Because efficient array manipulation will require low-level system access, the routines will actually be implemented using **external** clauses, introduced in a later chapter.

11.10 THE ADT CONNECTION

This section explores the implications of previous concepts. Some readers may prefer to skip to [AN ASSERTION TO AVOID A RESTRICTION*](#), [§11.11, page 379](#).

A class — you have heard this quite a few times by now — is an implementation of an abstract data type, whether formally specified or (as in many cases) just implicitly understood. As noted at the beginning of this chapter, we may view assertions as a way to re-introduce into the class the semantic properties of the underlying ADT. Let us perfect our understanding of assertion concepts by clarifying the connection of assertions to the components of an abstract data type specification.

Not just a collection of functions

As studied in the ADT chapter, an abstract data type is made of four elements:

- The name of the type, possibly with generic parameters (TYPES paragraph).
- The list of functions with their signatures (FUNCTIONS paragraph).
- The axioms (AXIOMS paragraph) expressing properties of the functions' results.
- The restrictions on the functions' applicability (PRECONDITIONS paragraph)

Simple-minded applications of abstract data types often overlook the last two parts. This removes much of the appeal of the approach, since preconditions and axioms express the semantic properties of the functions. If you omit them and simply view “stack” as encapsulating the (not specified further) operations *put*, *remove* etc., you retain the benefits of information hiding, but that is all. The notion of stack becomes an empty shell.

with no semantics other than suggested by the operation names. (And in the approach of this book that is of little comfort, since for reasons of structure, consistency and reusability we deliberately choose general names — *put*, *remove*, *item* ... — rather than concrete, type-specific names such as *push*, *pop* and *top*.)

This risk transposes to programming in an O-O language: the routines which are supposed to implement the operations of the corresponding abstract data types could in principle perform just about any operations. Assertions avert that risk by bringing the semantics back in.

Class features vs. ADT functions

To understand the relation between assertions and ADTs we need first to establish the relation between class features and their ADT counterparts — the ADT's functions. An earlier discussion introduced three categories of function: creators, queries and commands. As you will recall, the category of a function

$f: A \times B \times \dots \rightarrow X$

depended on where the ADT, say *T*, appeared among the types *A*, *B*, ... *X* involved in this signature:

- If *T* appears on the right only, *f* is a creator; in the class it yields a creation procedure.
- If *T* appears only on the left of the arrow, *f* is a query, providing access to properties of instances of the class. The corresponding features are either attributes or functions (collectively called queries, for classes as well as ADTs).
- If *T* appears on both the left and the right, *f* is a command function, which yields a new object from one or more existing objects. Often *f* will be expressed, at the implementation stage, by a procedure (also called a command) which modifies an object, rather than creating a new object as a function would do.

Expressing the axioms

From the correspondence between ADT functions and class features we can deduce the correspondence between semantic ADT properties and class assertions:

- A precondition for one of the specification's functions reappears as precondition clauses for the corresponding routine.
- An axiom involving a command function, possibly with one or more query functions, reappears as postcondition clauses of the corresponding procedure.
- Axioms involving only query functions reappear as postconditions of the corresponding functions or (especially if more than one function is involved, or if at least one of the queries is implemented as an attribute) as clauses of the invariant.
- Axioms involving constructor functions reappear in the postcondition of the corresponding creation procedure.

At this point you should go back to the preconditions and axioms of the ADT *STACK* and compare them with the assertions of class *STACK4* (including those of *STACK2*).

See “Function categories”, page 134.


See “Attributes and routines”, page 173.

Exercise E11.2, page 409. The ADT specification is on page 138.

Non-mathematical
readers may skip this
section

The abstraction function

It is instructive to think of the preceding observations in terms of the following figure, inspired by the discussion in [Hoare 1972a], which pictures the notion “ C is a correct implementation of A ”.



A is an abstract data type, and C as a class implementing it. For an abstract function af of the ADT specification — of which we assume for simplicity that it yields a result also of type A — there will be a concrete feature cf in the class.

The arrows labeled a represent the **abstraction function** which, for any instance of the class, or “concrete object”, yields the abstract object (instance of the ADT) that it represents. As will be seen, this function is usually partial, and the inverse relation is usually not a function.

The implementation is correct if (for all functions af applicable to abstract data types, and their implementations cf) the diagram is commutative, that is to say:

$$\text{Class-ADT Consistency property} \\ (cf; a) = (a; af)$$

where $;$ is the composition operator between functions; in other words, for any two functions f and g , $f; g$ is the function h such that $h(x) = g(f(x))$ for every applicable x . (The composition $f; g$ is also written $g \circ f$ with the order of the operands reversed.)

The property states that for every concrete object $CONC_1$, it does not matter in which order you apply the transformation (abstract af or concrete cf) and the abstraction; the two paths, represented by dotted lines, lead to the same abstract object $ABST_2$. The result is the same whether you:

- Apply the concrete transformation cf , then abstract the result, yielding $a(cf(CONC_1))$.
- Abstract first, then apply the abstract transformation af , yielding $af(a(CONC_1))$.

Implementation invariants

Certain assertions appear in invariants although they have no direct counterparts in the abstract data type specifications. These assertions involve attributes, including some secret attributes which, by definition, would be meaningless in the abstract data type. A simple example is the following properties appearing in the invariant of *STACK4*:

count *non_negative*: $0 \leq count$
count_bounded: $count \leq capacity$

Such assertions constitute the part of the class invariant known as the **implementation invariant**. They serve to express the consistency of the representation chosen in the class (here by attributes *count*, *capacity* and *representation*) vis-à-vis the corresponding abstract data type.

The figure on the previous page helps understand the concept of implementation invariant. It illustrates the characteristic properties of the abstraction function a (represented by the vertical arrows), which we should explore a little further.

First, is it correct to talk about a as being the abstraction function, as suggested by the upwards arrows representing a in the preceding figure? Recall that a function (partial or total) maps every source element to at most one target element, as opposed to the more general case of a relation which has no such restriction. If we go downwards rather than upwards in the figure and examine the inverse of a , which we may call the **representation relation**, we will usually find it not to be a function, since there are in general many possible representations of a given abstract object. In the array implementation that represents every stack as a pair $\langle representation, count \rangle$, an abstract stack has many different representations, as illustrated by the figure on the facing page; they all have the same value for *count* and for the entries of array *representation* between indices l and *count*, but the size *capacity* of the array can be any value greater than or equal to *count*, and the array positions beyond index *count* may contain arbitrary values.

Since the class interface is restricted to the features directly deduced from the ADT's functions, clients have no way of distinguishing between the behaviors of several concrete objects that all represent the same abstract object (that is to say, all have the same a value). Note in particular that procedure *remove* in *STACK4* does its job simply by executing


count := *count* – 1

without bothering to clear the previous top entry, now at index *count* + 1; changing an entry of index higher than *count* modifies a concrete stack object *CS*, but has no effect on the associated abstract stack $a(CS)$.

So the implementation relation is usually not a function. But its inverse the abstraction function a (the upwards arrows in both figures) is indeed a function since every concrete object represents at most one abstract object. In the stack example, every valid $\langle representation, count \rangle$ pair represents just one abstract stack (the stack with *count* elements, given, from the bottom up, by the entries of *representation* at indices l to *count*).

STACK4 and its
invariant appeared
on page 365.

Same abstract object, two representations



Both of the concrete stacks in this figure are implementations of the abstract stack consisting of three elements of values 342, -133 and 5 from the bottom up. That a is a function is a universal requirement: if the same concrete object could be interpreted as implementing more than one abstract object, the chosen representation would be ambiguous and hence inadequate. So it is proper that the arrow associated with a points up in all the figures depicting connections between abstract and concrete types. (The discussion for inheritance will suggest a similar convention.)

The abstraction function a is usually a *partial* function: not every possible concrete object is a valid representation of an abstract object. In the example, not every $\langle \text{representation}, \text{count} \rangle$ pair is a valid representation of an abstract stack: if representation is an array of capacity three and count has value 4, they do not together represent a stack. Valid representations (members of the domain of the abstraction function) are those pairs for which count has a value between zero and the size of the array. This property is the implementation invariant.

In mathematical terms, the implementation invariant is the characteristic function of the **domain** of the abstraction function, that is to say, the property that defines when that function is applicable. (The characteristic function of a subset A is the boolean property that is true on A and false elsewhere.)

The implementation invariant is the one part of the class's assertions that has no counterpart in the abstract data type specification. It relates not to the abstract data type, but to its representation. It defines when a candidate concrete object is indeed the implementation of one (and then only one) abstract object.

11.11 AN ASSERTION INSTRUCTION

The uses of assertions seen so far — preconditions, postconditions and class invariants — are central components of the method. They establish the connection between object-oriented software construction and the underlying theory (abstract data types). Class invariants, in particular, cannot be understood, or even discussed, in a non-O-O approach.

Some other uses of assertions, although less specific to the method, are also precursors in a systematic software development process and should be part of our notation. They include the **check** instruction, as well as loop correctness constructs (loop invariant and variant) which will be reviewed in the next section.

The **check** instruction serves to express the software writer's conviction that a certain property will be satisfied at certain stages of the computation. Syntactically, the construct is an instruction, written under the form

```

check
  assertion1 clause1
  assertion2 clause2
  ...
  assertionn clausen
end
```

Including this instruction in the text of a routine is a way to state that:

“Whenever control reaches this instruction at execution time, the assertion shown (as given by its assertion clauses) will hold.”

This is a way to reassure yourself that certain properties are satisfied, and (even more importantly) to make explicit for future readers of your software the hypotheses on which you have been relying. Writing software requires making frequent assumptions about properties of the objects of your system; as a trivial but typical example, any function call of the form $\text{sqrt}(x)$, where sqrt is a routine requiring a non-negative argument, relies on the assumption that x is positive or zero. This assumption may be immediately obvious from the context, for example if the call is part of a conditional instruction of the form

```
if  $x >= 0$  then  $y := \text{sqrt}(x)$  end
```

but the justification may also be more indirect, based for example on an earlier instruction that computed x as the sum of two squares:

```

 $x := a^2 + b^2$ 
...
 $x := a^2 + b^2$ 
... Other instructions ...
check
   $x >= 0$ 
end
-- Because  $x$  was computed above as a sum of squares.
```

Not if ... **then** ... protects the call to *sqrt* in this example; the **check** indicates that the call is correct. It is good practice to include, as here, a comment stating the reason invoked to support the assumption (“*Because x...*”). The extra two steps of indentation for the instruction are also part of the recommended style; they suggest that the instruction is not meant, in normal circumstances, to affect the algorithmic progression of the routine.

This example is typical of what is probably the most useful application of the **check** instruction: adding such an instruction just before a call to a routine that has a certain precondition (here we may assume that *sqrt* has a precondition requiring its argument to be non-negative), when you are convinced that the call satisfies the precondition but this is not immediately obvious from the context. As another example assume *s* is a stack and you include in your code a call

```
s.remove
```

at a position where you are certain that *s* is not empty, for example because the call has been preceded by *n* “**put**” and *m* “**remove**” instructions with $n > m$. Then there is no need to protect the call by an **if not s.empty then...**; but if the reason for the correctness of the call is not immediately obvious from the context, you may want to remind the reader that the omission of any protection was a conscious decision, not an oversight. You can achieve this by adding before the call the instruction

```
check not s.empty end
```

A variant of this case occurs when you write a call of the form *x.f* with the certainty that *x* is not void, so that you do not need to enclose this call in a conditional instruction **if x != Void then ...**; but the non-vacuity argument is not obvious from the context. We encountered this in the procedures *put* and *remove* of our “protected stack” class *STACK3*. The body of *put* used a call to the corresponding procedure in *STACK2*, as follows:

```
This is from the body
of put on page 382.

if full then
    error := Overflow
else
    check representation != Void end
    representation.put(x); error := 0
end
```

Exercise E11.5, page 410, asks you for the implementation invariant of

```
Here a reader might think the call representation.put(x) in the else potentially
unsafe since it is not preceded by a test for representation != Void. But if you examine the
class text you will realize that if full is false then capacity must be positive and hence
representation cannot be void. This is an important and not quite trivial property, which
should be part of the implementation invariant of the class. In fact, with a fully stated
implementation invariant, we should rewrite the check instruction as:
```

```
check
    representation_exists; representation != Void
    – Because of clause representation_exists, if not_full of the
    – implementation invariant.
end
```

In ordinary approaches to software construction, although calls and other operations often (as in the various preceding examples) rely for their correctness on various assumptions, these assumptions remain largely implicit. The developer will convince himself that a certain property always holds at a certain point, and will put this analysis to good use in writing the software text; but after a while all that survives is the text; the rationale is gone. Someone — even the original author, a few months later — who needs to understand the software, perhaps to modify it, will not have access to the assumption and will have to figure out from scratch what in the world the author may have had in mind. The **check** instruction helps avoid this problem by encouraging you to document your non-trivial assumptions.

As with the other assertion mechanisms of this chapter, the benefit goes beyond helping you get things *right* in the first place, to helping you find that you got them *wrong*. You can, using a compilation option, turn the **check** into a true executable instruction, which will do nothing if all its assertion clauses are true, but will produce an exception and stop execution if any of them is false. So if one of your assumptions was actually not justified you should find out quickly. The mechanisms for enabling **check-checking** will be reviewed shortly.

11.12 LOOP INVARIANTS AND VARIANTS

Our last assertion constructs help us get loops right. They nicely complement the mechanisms seen so far, but are not really specific to the object-oriented method, so it is all right to skip this section on first reading.

Loop trouble

The ability to repeat a certain computation an arbitrary number of times without succumbing to exhaustion, indeed without experiencing any degradation whatsoever, is the principal difference between the computational abilities of computers and those of humans. This is why loops are so important; just imagine what you could do in a language that only has the other two principal control structures, sequencing and conditional instructions, but no loops (and no support for recursive routine calls, the other basic mechanism permitting iterative computations).

But with power comes risk. Loops are notoriously hard to get right. Typical trouble includes:

- “Off-by-one” errors (performing one iteration too many or too few).
- Improper handling of borderline cases such as empty structures; for example a loop may work properly on a large array, but fail when the array has zero or one element.
- Failure to terminate (“infinite looping”) in some cases.

Binary search — a staple of Computing Science 101 courses — is a good illustration of how tricky loops can be even when they appear trivial. Consider an array *r* of integers assumed to be in increasing order and indexed from 1 to *n*; binary search is a way to decide whether a certain integer value *x* appears in the array; if the array has no elements, the

If skipping, go to “USING ASSERTIONS”, §11.13, page 390.

See exercise E11.7, page 410.

answer is no; if the array has one element, the answer is yes if and only if that element has value x ; otherwise compare x to the element at the array's middle position, and repeat on the lower or higher half depending on whether that element is greater or lesser than x . The four loop algorithms below all attempt to implement this simple idea; unfortunately all are wrong, as you are invited to check by yourself by finding, for each of them, a case in which it will not work properly.

Recall that $t @ m$ denotes the element at index t in array t . The $//$ operator denotes integer division, for example $7 // 2$ and $6 // 2$ have value 3. The loop syntax is explained next but should be self-explanatory; the **from** clause introduces the loop initialization.

BS1	BS2
<pre>from i := 1; j := n until i = j loop m := (i + j) // 2 if t @ m <= x then i := m else j := m end end Result := (x = t @ i)</pre>	<pre>from i := 1; j := n; found := false until i = j and not found loop m := (i + j) // 2 if t @ m < x then i := m + 1 elseif t @ m = x then found := true else j := m - 1 end end Result := found</pre>
BS3	BS4
<pre>from i := 0; j := n until i = j loop m := (i + j + 1) // 2 if t @ m <= x then i := m + 1 else j := m end end if i >= 1 and i <= n then Result := (x = t @ i) else Result := false end</pre>	<pre>from i := 0; j := n + 1 until i = j loop m := (i + j) // 2 if t @ m <= x then i := m + 1 else j := m end end if i >= 1 and i <= n then Result := (x = t @ i) else Result := false end</pre>

Getting loops right

The judicious use of assertions can help avoid such problems. A loop may have an associated assertion, the *loop invariant* (not to be confused with the class invariant for the enclosing class): it may also have a *loop variant*, not an assertion but an integer expression. The invariant and variant will help us guarantee that a loop is correct.

To understand these notions it is necessary to realize that a loop is always a way to compute a certain result by **successive approximations**.


Take the trivial example of computing the maximum value of an array of integers using the obvious algorithm:

```
maxarray (t: ARRAY [INTEGER]; INTEGER is
  -- The highest of the values in the entries of t
  require
    t.capacity >= 1
  local
    i: INTEGER
  do
    from
      i := t.lower
      Result := t @ lower
    until i = t.upper loop
      i := i + 1
      Result := Result.max (t @ i)
    end
  end
end
```

We initialize i to the array's lower bound $i := t.lower$ and the entry *Result* representing the future result to the value of the associated entry $t @ lower$. (We know that this entry exists thanks to the routine's precondition, which states that the array has at least one element.) Then we iterate until i has reached the upper bound, at each stage increasing i by one and replacing *Result* by the value of $t @ i$, the element at index i , if higher than the previous value of *Result*. (We rely on a *max* function for integers: $a.max(b)$, for two integers a and b , is the maximum of their values.)

This computation works by successive approximations. We approach the array by its successive slices: $[lower, lower]$, $[lower, lower+1]$, $[lower, lower+2]$ and so on up to the full approximation $[lower, upper]$.


Approximating
an array by
successive slices



The invariant property is that at each stage through the loop *Result* is the maximum of the current approximation of the array. This is true after the initialization, since the instructions in the **from** clause ensure that *Result* is the maximum of the first approximation, the trivial slice $[lower, lower]$ consisting of just one element. Then on each iteration we extend the slice by one element — improving our approximation of the array — and make sure to maintain the invariant by updating *Result* if the new value is higher than the previous maximum. At the end, the approximation covers the entire array, and since we have maintained invariant the property that *Result* is the maximum of the current approximation we know that it now is the maximum of the array as a whole.

Ingredients for a provably correct loop

The simple example of computing an array's maximum illustrates the general scheme of loop computation, which applies to the following standard situation. You have determined that the solution to a certain problem is an element belonging to an n -dimensional surface *POST*: to solve the problem is to find an element of *POST*. In some cases *POST* has just one element — *the* solution — but in general there may be more than one acceptable solution. Loops are useful when you have no way of shooting straight at *POST* but you see an indirect strategy: aiming first into an m -dimensional surface *INV* that includes *POST* (for $m > n$), then approaching *POST*, iteration by iteration, without ever leaving *INV*. The following figure illustrates this process.



A loop
computation
(from [M 1990])

A loop computation has the following ingredients:

- A goal *pos*, the postcondition, defined as a property that any satisfactory end state of the computation must satisfy. Example: “*Result* is the maximum value in the array”. The goal is represented in the figure by the set of states *POST* satisfying *post*.
- An invariant property *inv*, which is a generalization of the goal, that is to say includes the goal as a special case. Example: “*Result* is the maximum value in a non-empty array slice beginning at the lower bound”. The invariant is represented in the figure by the set of states *INV* satisfying *inv*.
- An initial point *init* which is known to be in *INV*, that is to say to satisfy the invariant. Example: the state in which the value of *i* is the array's lower bound and the value of *Result* is that of the array element at that index, satisfying the invariant since the maximum of a one-element slice is the value of the element.
- A transformation *body* which, starting from a point in *INV* but not in *POST*, yields a point closer to *POST* and still in *INV*. In the example this transformation extends the array slice by one element, and replaces *Result* by the value of that element if higher than the previous *Result*. The loop body in function *maxarray* is an implementation of that transformation.
- An upper bound on the number of applications of *body* necessary to bring a point in *INV* to *POST*. This will be the variant, as explained next.

Computations by successive approximations are a mainstay of numerical analysis, but the idea applies more broadly. An important difference is that in pure mathematics we accept that a series of approximations may have a limit even though it cannot reach it through a finite number of approximations: the sequence $1, 1/2, 1/3, 1/4, \dots, 1/n, \dots$ has

limit 0 but no element of the sequence has value zero. In computing, we want to see the results on our screen during our lifetime, so we insist that all approximation sequences reach their goal after a finite number of iterations.

Computer implementations of numerical algorithms also require finite convergence: even when the mathematical algorithm would only converge at infinity, we cut off the approximation process when we feel that we are close enough.

The practical way to guarantee termination of a loop process is to associate with the loop an integer quantity, the loop variant, which enjoys the following properties:

- The variant is always non-negative.
- Any execution of the loop body (the transformation called *body* in the figure) decreases the variant.

Since a non-negative integer quantity cannot decrease forever, your ability to exhibit such a variant for one of your loops guarantees that the loop will always terminate. The variant is an upper bound, for each point in the sequence, of the maximum number of applications of *body* that will land the point in *POST*. In the array maximum computation, a variant is easy to find: $t.upper - i$. This satisfies both conditions:

- Because the routine precondition requires $t.capacity$ to be positive (that is to say, the routine is only applicable to non-empty arrays) and the invariant of class *ARRAY* indicates that $capacity = upper - lower + 1$, the property $i \leq t.upper$ (part of the loop's invariant) will always be satisfied when i is initialized to $t.lower$.
- Any execution of the loop body performs the instruction $i := i + 1$, reducing the variant by one.

In this example the loop is simply an iteration over a sequence of integer values in a finite interval, known in common programming languages as a “for loop” or a “DO loop”; termination is not difficult to prove, although one must always check the details (here, for example, that i always starts no greater than $t.upper$ because of the routine's precondition). For more sophisticated loops, the number of iterations is not that easy to determine in advance, so ascertaining termination is more of a challenge: the only universal technique is to find a variant.

One more notion is needed to transform the scheme just outlined into a software text describing a loop: we need a simple way of determining whether a certain iteration has reached the goal (the postcondition) *post*. Because the iteration is constrained to remain within *INV*, and *POST* is part of *INV*, it is usually possible to find a condition *exit* such that an element of *INV* belongs to *POST* if and only if it satisfies *exit*. In other words, the postcondition *post* and the invariant *inv* are related by

$$post = inv \text{ and } exit$$

so that we can stop the loop — whose intermediate states, by construction, always satisfy *inv* — as soon as *exit* is satisfied. In the *maxarray* example, the obvious *exit* condition is $i = t.upper$: if this property is true together with the invariant, which states that *Result* is the maximum value in the array slice [$t.lower$, i], then *Result* is the maximum value in the array slice [$t.lower$, $t.upper$], hence in the entire array — the desired postcondition.

The invariant of class *ARRAY* appeared on page 374.

Loop syntax

The syntax for loops follows directly from the preceding rationale. It will include the elements listed as necessary:

- A loop invariant *inv* — an assertion.
- An exit condition *exit*, whose conjunction with *inv* achieves the desired goal.
- A variant *var* — an integer expression.
- A set of initialization instructions *init*, which always produces a state that satisfies *inv* and makes *var* non-negative.
- A set of body instructions *body* which, when started in a state where *inv* holds and *var* is non-negative, preserves the invariant and decreases the variant while keeping it non-negative (so that the resulting state still satisfies *inv* and has for *var* a value that is less than before but has not gone below zero).

The loop syntax combining these ingredients is straightforward:

```

from
  init
  invariant
    inv
  variant
    var
  until
    exit
  loop
    body
  end
```

The *invariant* and *variant* clauses are optional. The **from** clause is required (but the *init* instructions may be empty). The effect of this instruction is to execute the *init* instructions and then, zero or more times, the *body* instructions; the latter are executed only as long as *exit* is false.

In Pascal, C etc. the loop would be a “while” loop, since the loop body is executed zero or more times, unlike the “**repeat ... until**” loop for which the body is always executed at least once. Here the test is an exit condition, not a continuation condition, and the loop syntax includes room for initialization. So the equivalent in Pascal of **from *init* until *exit* loop *body* end** is

```

init;
while not exit do body
```

Warning: this is Pascal, not the O notation.

With a variant and an invariant the loop for *maxarray* appears as

```

from
  i := t.lower; Result := t @ lower
invariant
  -- Result is the maximum of the elements of t at indices t.lower to i.
variant
  t.lower - i
until
  i = t.upper
loop
  i := i + 1
  Result := Result.max (t @ i)
end

```

Note that the invariant is expressed informally as a comment; the discussion section of this chapter will explain this limitation of the assertion language.

Here is another example, first shown without variant or invariant. The purpose of the following function is to compute the greatest common divisor (gcd) of two positive integers *a* and *b* with Euclid's algorithm:

```

gcd (a, b): INTEGER; INTEGER is
  -- Greatest common divisor of a and b
require
  a > 0; b > 0
local
  x, y: INTEGER
do
  from
    x := a; y := b
  until
    x = y
  loop
    if x > y then x := x - y else y := y - x end
  end
  Result := x
ensure
  -- Result is the greatest common divisor of a and b
end

```

How do we know that function *gcd* ensures its postcondition — that it indeed computes the greatest common divisor of *a* and *b*? One way to check this is to note that the following property is true after loop initialization and preserved by every iteration:

$x > 0; y > 0$

-- The pair $\langle x, y \rangle$ has the same greatest common divisor as the pair $\langle a, b \rangle$

This will serve as our loop invariant *inv*. Clearly, *INV* is satisfied after execution of the **from** clause. Also, if *inv* is satisfied before an execution of the loop body

if *x* > *y* **then** *x* := *x* - *y* **else** *y* := *y* - *x* **end**

under the loop continuation condition $x \neq y$, then *inv* will still be satisfied after execution of this instruction; this is because replacing the greater of two positive non-equal numbers by their difference leaves them positive and does not change their gcd.

We have shown *inv* to be satisfied before the first iteration and preserved by every iteration. It follows that on loop exit, when $x = y$ becomes true, *inv* still holds; that is to say:

$x = y$ and "The pair $\langle x, y \rangle$ has the same greatest common divisor as the pair $\langle a, b \rangle$ "

which implies that the gcd is *x* because of the mathematical property that the gcd of any integer *x* and itself is *x*.

How do we know that the loop will always terminate? We need a variant. If *x* is greater than *y*, the loop body replaces *x* by *x* - *y*; if *y* is greater than *x*, it replaces *y* by *y* - *x*. We cannot choose *x* as a variant, because we cannot be sure that an arbitrary loop iteration will decrease *x*; nor can we be sure that it will decrease *y*, so *y* is also not an appropriate variant. But we can be sure that it will decrease either *x* or *y*, and hence their maximum *x.max(y)*; this maximum will never become negative, so it provides the sought variant. We may now write the loop with all its clauses:

```

from
  x := a; y := b
invariant
  x > 0; y > 0
  -- The pair  $\langle x, y \rangle$  has the same greatest common divisor as the pair  $\langle a, b \rangle$ 
variant
  x.max(y)
until
  x = y
loop
  if x > y then x := x - y else y := y - x end
end

```

As noted, the **invariant** and **variant** clauses in loops are optional. When present, they help clarify the purpose of a loop and check its correctness. Any non-trivial loop may be characterized by an interesting invariant and variant; many of the examples in subsequent chapters include variants and invariants, providing insights into the underlying algorithms and their correctness.

11.1.3 USING ASSERTIONS

We have now seen all the constructs involving assertions and should review all the benefits that we can derive from them. There are four main applications:

- Help in writing correct software.
- Documentation aid.
- Support for testing, debugging and quality assurance.
- Support for software fault tolerance.

Only the last two assume the ability to monitor assertions at run time.

Assertions as a tool for writing correct software

The first use is purely methodological and perhaps the most important. It has been explored in detail in the preceding sections: spelling out the exact requirements on each routine, and the global properties of classes and loops, helps developers produce software that is correct the first time around, as opposed to the more common approach of trying to debug software into correctness. The benefits of precise specifications and a systematic approach to program construction cannot be overemphasized. Throughout this book, whenever we encounter a program element, we shall seek to express as precisely as possible the formal properties of that element.

The key idea runs through this chapter: the principle of **Design by Contract**. To use features from a certain module is to contract out for services. Good contracts are those which exactly specify the rights and obligations of each party, and the *limits* to these rights and obligations. In software design, where correctness and robustness are so important, we need to spell out the terms of the contracts as a prerequisite to enforcing them. Assertions provide the means to state precisely what is expected from and guaranteed to each side in these arrangements.

Using assertions for documentation: the short form of a class

The second use is essential in the production of reusable software elements and, more generally, in organizing the interfaces of modules in large software systems. Preconditions, postconditions and class invariants provide potential clients of a module with basic information about the services offered by the module, expressed in a concise and precise form. No amount of verbose documentation can replace a set of carefully expressed assertions, appearing *in the software itself*.

To learn how a particular project ignored this rule and lost an entire space mission at a cost of \$500 million, see the very last section of this chapter.

The automatic documentation tool **short** uses assertions as an important component in extracting from a class the information that is relevant to potential clients. The short form of a class is a high-level view of the class. It only includes the information that is useful to authors of client classes: so it does not show anything about secret features and, for public features, it does not show the implementation (the **do** clauses). But the short form does retain the assertions, which provide essential documentation by stating the contracts that the class offers to its clients.

*"POSTSCRIPT:
THE AFTERMATH
OF ROSA", page 411.*

*Chapter 23 discusses
short in detail.*

Here is the short form of class `STACK4`:

```
indexing
description: "Stacks: Dispenser structures with a Last-In, First-Out %
    %access policy, and a fixed maximum capacity"
class interface STACK4 [G] creation
make
feature -- Initialization
make (n: INTEGER) is
    -- Allocate stack for a maximum of n elements
require
non_negative_capacity: n >= 0
ensure
capacity_set: capacity = n
end
feature -- Access
capacity: INTEGER
    -- Maximum number of stack elements
count: INTEGER
    -- Number of stack elements
item: C is
    -- Top element
require
not_empty: not empty -- i.e. count > 0
end
feature -- Status report
empty: BOOLEAN is
    -- Is stack empty?
ensure
empty_definition: Result = (count = 0)
end
full: BOOLEAN is
    -- Is stack full?
ensure
full_definition: Result = (count = capacity)
end
```

*STACK4 appeared
on page 385, listed
on STACK2 from
page 350.*

```

feature -- Element change
  put (x: G) is
    -- Add x on top
    require
      not_full: not full
    ensure
      not_empty: not empty
      added_to_top: item = x
      one_more_item: count = old count + 1
    end
  remove is
    -- Remove top element
    require
      not_empty: not empty -- i.e. count > 0
    ensure
      not_full: not full
      one_fewer: count = old count - 1
    end
  invariant
    count_non_negative: 0 <= count
    count_bounded: count <= capacity
    empty_if_no_elements: empty = (count = 0)
  end -- class interface STACK4

```

This short form is not a syntactically valid class text (hence the use of **class interface** rather than the usual **class** to avoid any confusion), although it is easy to turn it into a valid *deferred* class, a notion to be seen in detail in our study of inheritance.

See chapter 36 about the environment.

In the ISE environment, you obtain the short form of a class by clicking on the corresponding button in a Class Tool displaying a class; you can generate plain text, as well as versions formatted for a whole host of formats such as HTML (for Web browsing), RTF (Microsoft's Rich Text Format), FrameMaker's MML, TeX, troff and others. You can also define your own format, for example if you are using some text processing tool with its specific conventions for specifying fonts and layout.

If you compare the short form's assertions to those of the class, you will notice that all the clauses involving *representation* have disappeared, since that attribute is not exported.

The short form of documentation is particularly interesting for several reasons:

- The documentation is at a higher level of abstraction than what it describes, an essential requirement for quality documentation. The actual implementation, describing the *how*, has been removed, but the assertions, explaining the *what* (or in some cases the *why*) are still there. Note that the header comments of routines, which complement assertions by giving a less formal explanation of each routine's purpose, are retained, as well as the *description* entry of the **indexing** clause.

- A direct consequence of the Self-Documentation principle studied in our review of modularity concepts, the short form treats documentation not as a separate product but as information contained in the software itself. This means that there is only one product to maintain, a requirement that runs through this book. There is also, as a result, a much better chance that the documentation will be correct, since by having everything at the same place you decrease the risk of forgetting to update the documentation after a change to the software, or conversely.

- The short form can be extracted from the class by automatic tools. So the documentation is not something that you have to write; instead it is something that you ask "the computer" to produce, at the click of a mouse button, when you need it.

It is interesting to compare this approach with the notion of package interface present in Ada ("specification part"), where you write a module (package) in two parts: the interface and the implementation. Java uses a similar mechanism. The interface of a package has some similarities to the short form of a class, but also significant differences:

- There are no assertions, so all the "specification" that you can give is in the form of type declarations and comments.
- The interface is not produced by a tool but written separately. So the developer has to state many things twice: the headers of routines, their signatures, any header comments, declarations of public variables. This forced redundancy is tedious (it would be even more so with assertions) and, as always, raises the risk of inconsistency, as you may change one of the two parts and forget to update the other.

The short form (complemented by its variant the flat-short form, which deals with inheritance and is studied in a later chapter) is a principal contribution of the object-oriented method. In the daily practice of O-O development it appears all the time not just as a tool for documenting software, particularly reusable libraries, but also as the standard format in which developers and managers study existing designs, prepare new designs, and discuss proposed designs.

The reason for the central role of the short form in O-O development is that it finally fulfills the goal defined by the analysis of reusability requirements at the beginning of this book. There we arrived at the requirement for *abstracted modules* as the basic unit of reuse. A class in its short (or flat-short) form is the abstracted module that we have been seeking.

Monitoring assertions at run time

It is time now to deal in full with the question "what is the effect of assertions at run time?". As previewed at the beginning of this chapter, the answer is up to the developer, based on a compilation option. To set that option, you should not, of course, have to change the actual class texts; you will rely instead on the Ace file. Recall that an Ace file, written in Lace, allows you to describe how to assemble and compile a system.

Recall too that Lace is just one possible control language for assembling O-O systems, not an immutable component of the method. You will need something like Lace, even if it is not exactly Lace, to go from individual software components to complete compilable systems.

"Self-Documentation", page 54.

"The flat-short form", page 543.

"Raise of abstracted modules", page 73.

"Lace and Ace files were introduced in 'Assembling a system', page 198.

Here is how to adapt a simple Ace (the one used as example in the original presentation of Lace) to set some assertion-monitoring options:

Warning: this text is in Lace, not in the O-O notation.

```
system painting root
  GRAPHICS
default
  assertion (require)
cluster
  base_library: "library\base"
  graphical_library: "library\graphics"
option
  assertion (all): BUTTON, COLOR_BITMAP
end
  painting_application: "user\application"
option
  assertion (no)
end
end -- system painting
```

The **default** clause indicates that for most classes of the system only preconditions will be checked (**require**). Two clusters override this default: *graphical_library*, which will monitor all assertions (**all**), but only for classes *BUTTON* and *COLOR_BITMAP*; and *painting_application*, which has disabled any assertion checking for all its classes. This illustrates how to define an assertion monitoring level for the system as a whole, for all the classes of a cluster, or for some classes only.

The following assertion checking levels may appear between parentheses in *assertion (...)*:

- **no**: do not execute anything for assertions. In this mode assertions have no more effect on execution than comments.
- **require**: check that preconditions hold on routine entry.
- **ensure**: check that postconditions hold on routine exit.
- **invariant**: check that class invariants hold on routine entry and exit for qualified calls.
- **loop**: check that loops invariants hold before and after every loop iteration, and that variants decrease while remaining non-negative.
- **check**: execute *check* instructions by checking that the corresponding assertions hold. **all** is a synonym for **check**.

Excluding **no**, each of these levels implies the previous ones; in particular it does not make sense to monitor postconditions unless you also monitor preconditions, since the principles of Design by Contract indicate that a routine is required to ensure its postcondition only if it was called with its precondition satisfied (otherwise “the customer is wrong”). This explains why **check** and **all** are synonyms.

A qualified call is a call from the outside, or from an object, to a method of a class. See “Qualified and unqualified calls”, page 186.

If turned on, assertion monitoring will have no visible effect, except for the CPU cycles that it takes away from your computation, as long as the assertions that it monitors all evaluate to true. But having any assertion evaluate to false is a rather serious event which will usually lead to termination. Actually it will trigger an exception, but unless you have taken special measures to catch the exception (see next) everything will stop. An *exception history table* will be produced, of the general form

```
Failure: object: O2 class: YOUR_CLASS routine: your_routine
Cause: precondition violation, clause: no_too_small
Called by: object: O2 class: YOUR_CLASS routine: his_routine
Called by: object: O1 class: HER_CLASS routine: her_routine
...
```

This gives the call chain, starting from the routine that caused the exception, the object to which it was applied and its generating class. Objects are identified by internal codes. The form shown here is only a sketch; the discussion of exceptions will give a more complete example of the exception history table.

The optional labels that you can add to the individual clauses of an assertion, such as

```
no_too_small in
your_routine (x: INTEGER) is
  require
    no_too_small: x >= Minimum_value
```

...
prove convenient here, since they will be listed in the exception trace, helping you identify what exactly went wrong.

How much assertion monitoring?

What level of assertion tracing should you enable? The answer is a tradeoff between the following considerations: how much you trust the correctness of your software; how crucial it is to get the utmost efficiency; how serious the consequences of an undetected run-time error can be.

In extreme cases, the situation is clear:

- When you are debugging a system, or more generally testing it prior to release, you should enable assertion monitoring at the highest level for the classes of the system (although not necessarily for the libraries that it uses, as explained next). This ability is one of the principal contributions to software development of the method presented in this book. Until they have actually had the experience of testing a large, assertion-loaded system using the assertion monitoring mechanisms described in this section, few people realize the power of these ideas and how profoundly they affect the practice of software development.
- If you have a fully trusted system in an efficiency-critical application area — the kind where every microsecond counts — you may consider removing all monitoring.

See page 421 for the detailed form.

The last advice is somewhat paradoxical since in the absence of formal proving techniques (see the discussion section of this chapter) it is seldom possible to “trust a system fully” — except by monitoring its assertions. This is a special case of a general observation made with his customary eloquence by C.A.R. Hoare:

From [Hoare 1973].

It is absurd to make elaborate security checks on debugging runs, when no trust is put in the results, and then remove them in production runs, when an erroneous result could be expensive or disastrous. What would we think of a sailing enthusiast who wears his life-jacket when training on dry land but takes it off as soon as he goes to sea?

An interesting possibility is the option that only checks preconditions: **assertion (require)**. In production runs — that is to say, past debugging and quality assurance — it has the advantage of avoiding catastrophes that would result from undetected calls to routines outside of their requirements, while costing significantly less in run-time overhead than options that also check postconditions and invariants. (Invariants, in particular, can be quite expensive to monitor since the method suggests writing rich invariants that include all relevant consistency conditions on a class, and the invariant is checked on entry and exit for every qualified call.)

Precondition checking is indeed the default compilation option if you do not include a specific **assertion** option in your Ace, so that the clause **default assertion (require)** appearing in the example Ace for system **painting** was not necessary.

Second Assertion Violation rule, page 347.

This option is particularly interesting for libraries. Remember the basic rule on assertion violations: a violated precondition indicates an error in the client; a violated postcondition or invariant indicates an error in the supplier. So if you are relying on reusable libraries that you assume to be of high quality, it is generally not desirable to monitor their postconditions and invariants; this would mean that you suspect the libraries themselves, and although the possibility of a library error is of course always open it should only be investigated (for a widely used library coming from a reputable source) once you have ruled out the presence, *a priori* much more likely, of an error in your own client software. But even for a perfect library it is useful to check **preconditions**: the goal is to find errors in client software.

See the class text starting on page 373.

Perhaps the most obvious example is array bound checking. In the **ARRAY** class we saw that **put**, **item** and the later’s synonym **infix “@”** all had the precondition clauses

```
index_not_100_small; lower <= i
index_not_100_large; i <= upper
```

Enabling precondition checking for the class solves a well-known problem of any software that uses arrays: the possibility of an out-of-bounds array access, which will usually scoop some memory area containing other data or code, causing ravages. Many compilers for conventional programming languages offer special compilation options to monitor array access at run time. But in object technology, just as we treat arrays through general notions of class and object rather than special constructs, we can handle array bound monitoring through the general mechanism for precondition checking. Just use a version of **ARRAY** compiled with **assertion (require)**.

Should bounds always be checked? Hoare thinks so:

In our Algol compiler every occurrence of every subscript of every array element was on every occasion checked at run time against the declared bounds. Many years later we asked our customers whether they wished us to provide an option to switch off these checks in the interest of efficiency in production runs.

From [Hoare 1981]; slightly abridged.

Unanimously they urged us not to — they already knew how frequently index errors occur on production runs where failure could be disastrous. I note with fear and horror that even today, language designers and users have not learned this lesson. In any respectable branch of engineering, failure to observe such elementary precautions would have long been against the law.

These comments should be discussed not just for arrays but for preconditions in general. If indeed “index errors frequently occur on production runs” this must be true of other precondition violations too.

One may defend a less extreme position. (Some might of course see here an attempt at self-preservation, coming from a “language designer” who has provided a way to turn off assertion checking, through Lace options such as **assertion (no)**, and presumably does not like being branded as acting “against the law”.) First, a company which delivers software in which precondition errors “frequently occur on production runs” probably has a problem with its software quality practices, which run-time assertion monitoring will not solve. Monitoring addresses the symptoms (*faults* in the terminology introduced earlier in this chapter), not the cause (defects and errors). True, assertion monitoring is in such a case beneficial to the software’s end-users: however unpleasant it is to have a system end its interruption with some message spouting insults about preconditions and other venomous beasts unknown to a layman, this is better than continuing operation and producing bad results. But in the long term a practice of always delivering systems with some level of assertion monitoring also has negative effects: it can encourage among developers, even unconsciously, a happy-go-lucky attitude towards correctness, justified by the knowledge that if an error remains it will be caught by the users through an assertion violation, reported to the supplier, and fixed for the following release. So can’t we stop testing right now and start shipping?

It is hard to give an absolute answer to the question “should we leave some assertion monitoring on?” without some knowledge of the performance overhead of assertion monitoring. If adding some monitoring multiplied the execution time by ten, few people outside of the mission-critical-computing community would support Hoare’s view; if the overhead were two percent, few people would disagree with it. In practice, of course, the penalty will be somewhere in-between.

How much is it, by the way? This clearly depends on what the software does and how many assertions it has, but it is possible to give empirical observations. In ISE’s experience the cost for monitoring preconditions (the default option, including of course array bounds checking) is on the order of 50%. What is frustrating is that more than a 75% of that cost is due not to precondition checking per se but to the supporting machinery of monitoring calls — recording every routine entry and every routine exit — so that if a precondition fails the environment can say which one and where. (A message of the form **Execution**

stopped because some assertion was violated somewhere would not be very useful.) This may be called the Precondition Checking Paradox: precondition checking is by itself cheap enough, but to get it you have to pay for something else. As to postcondition and invariant checking, they can bring the penalty to 100% to 200%. (Although circumstances vary, preconditions are often relatively simple consistency conditions such as $x > 0$ or $a \neq \text{Void}$, whereas many postconditions and invariants express more advanced semantic properties.)

One might fear that bringing performance into this discussion may lead to compromising on correctness, against the principle expressed at the beginning of this book:

Necessary as tradeoffs between quality factors may be, one factor stands out from the rest: correctness. There is never any justification for compromising on correctness for the sake of other concerns, such as efficiency. If the software does not perform its function, the rest is useless.

Considering performance when we decide whether to leave assertion monitoring on is not, however, a violation of this principle. The point is not to sacrifice correctness for efficiency, but to determine what we should do for systems that are *not* correct — obviously because we have not worked hard enough at making them correct.

In fact, efficiency may be part of correctness. Consider a meteorological system that takes twelve hours to predict the next-day's weather (two hours would be more useful, of course). The system has been thoroughly optimized; in particular it does not have run-time checking for array bound violations or other such faults. It has also undergone careful development and extensive testing. Now assume that adding the run-time checks multiplies the execution time by two, giving a forecasting system that takes 24 hours to predict tomorrow's weather. Would you enable these checks? No.

Although the examples that first come to mind when discussing such performance vs. safety issues tend to be of the Parlor-against-Scud variety, I prefer the weather forecasting example because here one cannot dismiss the efficiency issue offhand by saying "just buy a faster microprocessor". In meteorological computing, the hardware tends *directly* to be the fastest parallel computer available on the market.

Let us not stop here but ask the really hard questions. Assume the original running time of twelve hours was with checking *enabled*. Would you disable it to get a six-hour forecast? Now assume that you also have the option of applying the improved efficiency to keep the same running time but use a more accurate forecasting model (since you can afford more grid points); would you do it? I think that in either case, if offered "an option to switch off the checks in the interest of efficiency in production runs", almost everyone will say yes.

So in the end the choice of assertion monitoring level at production time is not as simple as Hoare's rule suggests. But a few precise and strict principles do hold:

- Remember that a software system should be made reliable *before* it begins operation. The key is to apply the reliability techniques described in the software engineering literature, including those which appear in this chapter and throughout this book.
- If you are a project manager, *never let the developers assume* that the production versions will have checks turned on. Make everyone accept that — especially for the

biggest production runs, those which by nature make the consequences of potential errors most frightening — all checks may be off.

- Make sure that during development assertion checking is always turned on at least at the precondition level.
- Perform extensive testing with *all* the checks enabled. Also turn all checks on as soon as any bug is encountered during development.
- For the standard production versions, decide whether to choose a no-check version or a protected version (usually at the precondition level) based on your assessment, from an engineering perspective, of the relative weight of the three factors cited at the beginning of this discussion: how much you trust the correctness of your software (meaning in part how hard you have worked at making it correct and convincing yourself and others that it is); how crucial it is to get the utmost efficiency; and how serious the consequences of an undetected run-time error can be.
- If you decide to go for a no-check version, also include in your delivery a version that checks at least for preconditions. That way, if the system starts exhibiting abnormal behavior against all your expectations, you can ask the users — those at least who have not been killed by the first erroneous production runs — to switch to the checking version, helping you find out quickly what is wrong.

Used in this way, run-time assertion monitoring provides a remarkable aid for quickly weeding out any errors that may have survived a systematic software construction process.

11.1.4 DISCUSSION

The assertion mechanism presented in this chapter raises some delicate issues, which we must now examine.

Why run-time monitoring?

Should we really have to check assertions at run time? After all we were able, using assertions, to give a theoretical definition of what it means for a class to be correct every creation procedure should ensure the invariant, and every routine body, when started in a state satisfying the precondition and the invariant, should maintain the invariant and ensure the postcondition. This means that we should simply *prove* the $m + n$ corresponding properties mathematically (for m creation procedures and n exported routines), and then do away with run-time assertion monitoring.

We should, but we cannot. Although mathematical program proving has been an active area of research for many years, and has enjoyed some successes, it is not possible today to prove the correctness of realistic software systems written in full-fledged programming languages.

We would also need a more extensive assertion language. The IFL sublanguage, discussed below, could be used as part of a multi-tier proof strategy.

See "WHEN IS A CLASS CORRECT?", 11.9, page 370.

Even if proof techniques and tools eventually become available, one may suspect that run-time checks will *not* go away, if only to cope with hard-to-predict events such as hardware faults, and to make up for possible bugs in the proof software itself — in other words to apply the well-known engineering technique of multiple independent checking.

The expressive power of assertions

As you may have noted, the assertion language that we have used is essentially the language of boolean expressions, extended with a few concepts such as *old*. As a result, we may find it too restrictive when we would like to include in our classes some of the properties that were easy to express in the mathematical notation for abstract data types.

The assertions for stack classes provide a good example of what we can and cannot say. We found that many of the preconditions and axioms from the original ADT specification of chapter 6 gave assertion clauses; for example the axiom

A4 • *not empty* (*put* (*s*, *x*))

gives the postcondition *not empty* in procedure *put*. But in some cases we do not have the immediate counterpart in the class. None of the postconditions for *remove* in the stack classes given so far includes anything to represent the axiom

A2 • *remove* (*put* (*s*, *x*)) = *s*

We can of course add an informal property to the postcondition by resorting to a comment:

```
remove is
    -- Remove top element
    require
        not_empty: not empty -- i.e. count > 0
    do
        count := count - 1
    ensure
        not_full: not full
        one_fewer: count = old count - 1
        LIFO_policy: -- item is the last element pushed (by put)
                    -- and not yet removed, if any.
end
```

Both on page 398.

Similar informal assertions, syntactically expressed as comments, appeared in the loop invariants for *maxarray* and *gcd*.

In such a case, two of the principal uses of assertions discussed earlier remain applicable at least in part: help in composing correct software, and help in documentation (an assertion clause that is syntactically a comment will appear in the short form). The other uses, in particular debugging and testing, assume the ability to evaluate assertions and do not apply any more.

It would be preferable to express all assertions formally. The best way to reach this goal is to extend the assertion language so that it can describe arbitrary properties; this requires the ability to describe complex mathematical objects such as sets, sequences, functions and relations, and including the full power of first-order predicate calculus, which allows quantified expressions (“for all” and “there exists”). Formal specification languages exist which provide at least part of this expressive power. The best known are Z, VDM, Larch and OBJ-2; both Z and VDM have had object-oriented extensions, such as Object-Z, in recent years, and the last two were close to O-O concepts already. The bibliographic notes to chapter 6 provide references.

Including a full specification language into the language of this book would have completely changed its nature. The language is meant to be simple, easy to learn, applicable to all areas of software construction, and implementable efficiently (with a final run-time performance similar to that of Fortran and C, and a fast compilation process).

Instead, the assertion mechanism is an engineering tradeoff: it includes enough formal elements to have a substantial effect on software quality; but stops at the point of diminishing return — the threshold beyond which the benefits of more formality might start being offset by the decrease of learnability, simplicity and efficiency.

Determining that threshold is clearly a matter of personal judgment. I have been surprised that, for the software community at large, the threshold has not moved since the first edition of this book. Our field needs more formality, but the profession has not realized it yet.

So for the time being, and probably for quite a while, assertions will remain boolean expressions extended with a few mechanisms such as the *old* expression in postconditions. The limitation is not as stringent as it seems at first, because boolean expressions can use *function* calls.

Including functions in assertions

A boolean expression is not restricted to using attributes or local entities. We have already used the possibility of calling *functions* in assertions: the precondition for *put* in our stack classes was *not full*, where *full* is the function

```
full: BOOLEAN is
    -- Is stack full?
    do
        Result := (count = capacity)
    ensure
        full_definition: Result = (count = capacity)
    end
```

This is our little assertion secret: we get out of the stranglehold of propositional calculus — basic boolean expressions involving attributes, local entities and boolean operators such as *and*, *or*, *not* — thanks to function routines, which give us the power to compute a boolean value in any way we like. (You should not be troubled by the presence

of a postcondition in *full* itself, as it does not create any harmful circularity. Details shortly.)

Using function routines is a way to obtain more abstract assertions. For example, some people may prefer replacing the precondition of the array operations, expressed earlier as

```
index_not_too_small: lower <= i
index_not_too_large: i <= upper
```

by a single clause of the form

```
index_in_bounds: correct_index ()
```

with the function definition

```
correct_index (i: INTEGER): BOOLEAN is
  -- Is i within the array bounds?
do
  Result := (i >= lower) and (i <= upper)
ensure
  definition: Result = ((i >= lower) and (i <= upper))
end
```

Another advantage of the use of functions in assertions is that it may provide a way to circumvent the limitations on expressive power arising from the absence of first-order predicate calculus mechanisms. The informal invariant of our *maxarray* loop

```
-- Result is the maximum of the elements of t at indices t.lower to i
```

may be expressed formally as

```
Result = (t.slice (lower, i)).max
```

assuming a function *slice* which yields the set of elements between two indices of an array, and a function *max* which yields the maximum element in a set.

This approach has been explored in [M 1995a] as a way to extend the power of the assertion mechanism, possibly leading to a fully formal development method (that is to say, to software that may be *proven* correct mathematically). Two central ideas in this investigation are the use of libraries in any large-scale proof process, so that one could prove real, large-scale systems in a multi-tier proof structure using conditional proofs, and the definition of a restricted language of a purely applicative nature — IFL, for Intermediate Functional Language — in which to express the functions used in assertions. IFL is a subset of the notation of this book, which excludes some imperative constructs such as arbitrary assignments.

The risk that such efforts try to address is clear: as soon as we permit functions in assertions, we introduce potentially imperative elements (routines) into the heretofore purely applicative world of assertions. Without functions, we had the clear and clean separation of roles emphasized in the earlier discussion: instructions prescribe, assertions describe. Now we open the gates of the applicative city to the imperative hordes.

Yet it is hard to resist the power of using functions, as the alternatives are not without their drawbacks either.

"The imperative and the applicative", page 352.

- Including a full specification sublanguage could, as noted, cause problems of ease of learning and efficiency.

- Perhaps worse, it is not even clear that commonly accepted assertion languages would suffice. Take what most people versed in these topics would suggest as the natural candidate: first-order predicate calculus. This formalism will not enable us to express some properties of immediate interest to developers and common use in assertions, such as "the graph has no cycles" (a typical invariant clause). Mathematically this would be stated as $r^+ \cap r = \emptyset$ where r is the graph's relation and $^+$ is transitive closure. Although it is possible to conceive of a specification language that supports these notions, most do not.

This is all the more troubling because, for a programmer, writing a boolean-valued function routine *cycle* that explores the graph and returns true if and only if there is a cycle, is not particularly hard. Such examples provide a strong argument for contenting ourselves with a basic assertion language and using functions for anything beyond its expressive power.

But the need to separate applicative and imperative elements remains. Any function routine used in an assertion to specify the properties of a software element should be "beyond reproach", more precisely beyond imperative reproach; it should not cause any permanent change of the abstract state.

This informal requirement is clear enough in practice; the IFL sublanguage formalizes it by excluding all the imperative elements which either change the global state of the system or do not have trivial applicative equivalents, in particular:

- Assignments to attributes.
- Assignments in loops.
- Calls to routines not themselves in IFL.

If you exert the proper care by sticking to functions that are simple and self-evidently correct, the use of function routines in assertions can provide you with a powerful means of abstraction.

A technical point may have caught your attention. A function f used by an assertion for a routine r (or the invariant of the class containing r) may itself have assertions, as illustrated by both the *full* and *correct_index* examples. This raises a potential problem for run-time assertion monitoring: if as part of a call to r we evaluate an assertion and this causes a call to f , we do not want the call to evaluate any assertion that f itself may have. For one thing, it is easy to construct examples that would cause infinite recursion. But even without that risk it would be just wrong to evaluate the assertions of f . This would mean that we treat as peers the routines of our computation, such as r , and their assertions's functions, such as f — contradicting the rule that assertions should be on a

The transitive closure of a relation is obtained by iterating it any number of times. For "simple" predecessor is the transitive closure of "parent".

higher plane than the software they protect, and their correctness crystal-clear. The rule is simple:

Assertion Evaluation rule

During the process of evaluating an assertion at run-time, routine calls shall be executed without any evaluation of the associated assertions.

If a call to f occurs as part of assertion checking for r , that is too late to ask whether f satisfies its assertions. The proper time for such a question is when you decide to use f in the assertions applicable to r .

We can use an analogy introduced earlier. Think of f as a security guard at the entrance of a nuclear plant, in charge of inspecting the credentials of visitors. There are requirements on guards too. But you will run the background check on a guard in advance; not while he is screening the day's visitors.

Class invariants and reference semantics

The object-oriented model developed so far includes two apparently unrelated aspects, both useful:

- The notion of class invariant, as developed in this chapter.
- A flexible run-time model which, for various reasons detailed in an earlier chapter (in particular the modeling needs of realistic systems), makes considerable use of references.

See chapter 8, in particular "DEALING WITH REFERENCES, BENEFITS AND DANGERS", 8.9, page 265.

Unfortunately these individually desirable properties cause trouble when put together.

The problem is, once again, dynamic aliasing, which prevents us from checking the correctness of a class on the basis of that class alone. We have seen that the correctness of a class means $m + n$ properties expressing that (if we concentrate on the invariant INV , ignoring preconditions and postconditions which play no direct role here):

- P1 • Every one of the m creation procedures produces an object that satisfies INV .
- P2 • Every one of the n exported routines preserves INV .

These two conditions seem sufficient to guarantee that INV is indeed invariant. The proof is apparently trivial: since INV will be satisfied initially, and preserved by every routine call, it should by induction be satisfied at all stable times.

This informal proof, however, is not valid in the presence of reference semantics and dynamic aliasing. The problem is that attributes of an object may be modified by an operation on another object. So even if all $o.r$ operations preserve INV on the object OA attached to o , some operation $b.s$ (for b attached to another object) may destroy INV for OA. So even with conditions P1 and P2 satisfied, INV may not be an invariant.


Here is a simple example. Assume classes A and B , each with an attribute whose type is the other's class:

```
class A ... feature forward: B ... end
class B ... feature backward: A ... end
```

We require that following the *forward* reference (if defined) from an instance of A and then the *backward* reference from the corresponding B will yield the original A . This may be expressed as an invariant property of A :

```
round_trip: (forward != Void) implies (forward.backward = Current)
```

Here is a situation involving instances of both classes and satisfying the invariant:



Invariant clauses of the *round_trip* form are not uncommon; think of *forward* in class *PERSON* denoting a person's residence, and *backward* in class *HOUSE* denoting a house's resident. Then *round_trip* states that the resident of any person's residence is that person, a simple consistency requirement. Another example is the linked implementation of trees, where the attributes of a tree node include references to its first child and to its parent, introducing the following *round_trip*-style property in the invariant:

```
(first_child != Void) implies (first_child.parent = Current)
```

Assume, however, that the invariant clause of B , if any, says nothing about the attribute *backward*. The following version of A appears consistent with the invariant:

```
class A feature
  forward: B
  attach (b1: B) is
    -- Link b1 to current object.
    do
      forward := b1
    -- Update b1's backward reference for consistency:
    if b1 != Void then
      b1.attach(Current)
    end
  end
end
invariant
  round_trip: (forward != Void) implies (forward.backward = Current)
end
```

The call *b1.attach* is meant to restore the invariant after an update of *forward*. Class *B* must provide its own *attach* procedure:

```

class B feature
  backward: B
  attach (a1: A) is
    -- Link a1 to current object.
  do
    backward := a1
  end
end

```


Class *A* appears to be correct: a procedure-less creation instruction ensures the invariant *round_trip* (since it initializes *forward* to a void reference), and its sole procedure will always preserve *round_trip*. But consider execution of the following:

```

a1: A; b1: B
...
!! a1; !! b1
a1.attach (b1)
b1.attach (Void)

```

Here is the situation after the last instruction:



The invariant is violated on OA! This object is now linked to OB, but OB is not linked to OA since its *backward* field is void. (A call to *b1.attach (...)* could also have linked OB to an instance of *A* other than OA, which would be equally incorrect.)

What happened? Dynamic aliasing has struck again. The proof of correctness of class *A* outlined above is valid: every operation of the form *a1.r*, where *a1* is a reference to object OA, will preserve *round_trip* since the corresponding features of *A* (there there is only one, *attach*) have been designed accordingly. But this is not sufficient to preserve the consistency of OA, since properties of OA may involve instances of other classes, such as *B* in the example, and the proof says nothing about the effect of these other classes' features on the invariant of *A*.

This problem is important enough to deserve a name: **Indirect Invariant Effect**. It may arise as soon as we allow dynamic aliasing, through which an operation may modify

an object even without involving any entity attached to it. But we have seen how much we need dynamic aliasing; and the *forward-backward* scheme, far from being just an academic example, is as noted a useful pattern for practical applications and libraries.

What can we do? The immediate answer involves the conventions for run-time monitoring of assertions. You may have wondered why the effect of enabling assertion monitoring at the *assertion (invariant)* level was described as

“Check that class invariants hold on routine entry and exit for qualified calls.”

Page 304.

Why both entry and exit? Without the Indirect Invariant Effect, it would suffice to check the invariant when exiting qualified calls. (It is also checked at the end of creation calls.) But now we have to be more careful, since between the termination of a call and the beginning of the next one on the same object, some call may have affected that object even though its target was another object.

A more satisfactory solution would be to obtain a statically enforceable validity rule, which would guarantee that whenever the invariant of a class *A* involves references to instances of a class *B*, the invariant of *B* includes a mirror clause. In our example we can avoid trouble by including in *B* an invariant clause *trip_round* mirroring *round_trip*:

trip_round: (*backward* != Void) implies (*backward.forward* = Current)

It may be possible to generalize this observation to a universal mirroring rule. Whether such a rule indeed exists, solving the Indirect Invariant Effect and removing the need for double run-time monitoring, requires further investigation.

More to come

We are not done with Design by Contract. Two important consequences of the principles remain to be studied:

- How they lead to a disciplined exception handling mechanism; this is the topic of the next chapter.
- How they combine with inheritance, allowing us to specify that any semantic constraints that apply to a class also apply to its descendants, and that semantic constraints on a feature apply to its eventual redeclarations; this will be part of our study of inheritance.

More generally, assertions and Design by Contract will accompany us throughout the rest of this book, enabling us to check, whenever we write software elements, that we know what we are doing.

11.15 KEY CONCEPTS INTRODUCED IN THIS CHAPTER

- Assertions are boolean expressions expressing the semantic properties of classes and reintroducing the axioms and preconditions of the corresponding abstract data types.
- Assertions are used in preconditions (requirements under which routines are applicable), postconditions (properties guaranteed on routine exit) and class

invariants (properties that characterize class instances over their lifetime). Other constructs that involve assertions are loop invariants and the **check** instruction.

- A precondition and a postcondition associated with a routine describe a contract between the class and its clients. The contract is only binding on the routine inasmuch as calls observe the precondition; the routine then guarantees the postcondition on return. The notion of contracting provides a powerful metaphor for the construction of correct software.
- The invariant of a class expresses the semantic constraints on instances of the class. The invariant is implicitly added to the precondition and the postcondition of every exported routine of the class.
- A class describes one possible representation of an abstract data type; the correspondence between the two is expressed by the abstraction function, which is usually partial. The inverse relation is in general not a function.
- An implementation invariant, part of the class invariant, expresses the correctness of the representation vis-à-vis the corresponding abstract data type.
- A loop may have a loop invariant, used to deduce properties of the result, and a variant, used to ascertain termination.
- If a class is equipped with assertions, it is possible to define formally what it means for the class to be correct.
- Assertions serve four purposes: aid in constructing correct programs; documentation aid; debugging aid; basis for an exception mechanism.
- The assertion language of our notation does not include first-order predicate calculus, but can express many higher-level properties through function calls, although the functions involved must be simple and of unimpeachable correctness.
- The combination of invariants and dynamic aliasing raises the Indirect Invariant Effect, which may cause an object to violate its invariant through no fault of its own.

11.16 BIBLIOGRAPHICAL NOTES

According to Tony Hoare:

From [Hoare 1981].

An early advocate of using assertions in programming was none other than Alan Turing himself. On 24 June 1950 at a conference in Cambridge, he gave a short talk entitled "Checking a Large Routine" which explains the idea with great clarity. "How can one check a large routine in the sense that it's right? In order that the man who checks may not have too difficult a task, the programmer should make a number of definite assertions which can be checked individually, and from which the correctness of the whole program easily follows."

The notion of assertion as presented in this chapter comes from the work on program correctness pioneered by Bob Floyd [Floyd 1967], Tony Hoare [Hoare 1969] and Edsger

Dijkstra [Dijkstra 1976], and further described in [Gries 1981]. The book *Introduction to the Theory of Programming Languages* [M 1990] presents a survey of the field.

The notion of class invariant comes from Hoare's work on data type invariants [Hoare 1972a]. See also applications to program design in [Jones 1980] [Jones 1986]. A formal theory of morphisms between abstract data types may be found in [Goguen 1978].

Formal specification languages include Z, VDM, OBJ-2 and Larch; see the bibliographical references to chapter 6. Object-oriented formal specification languages include Object Z, Z++, MoorZ, OOZE, SmallVDM and VDM++, all of which are described in [Lano 1994] which gives many more references.

The IEEE Computer Society publishes standards for the terminology of software errors, defects, failures [IEEE 1990] [IEEE 1993]. Its Web page is at <http://www.computer.org>.

Surprisingly, few programming languages have included syntactical provision for assertions; an early example (the first to my knowledge) was Hoare's and Wirth's Algol W [Hoare 1966], the immediate precursor of Pascal. Others include Alphasol [Shaw 1981] and Euclid [Lampson 1977], which were specifically designed to allow the construction of provably correct programs. The connection with object-oriented development introduced by the notation developed in this book was foreshadowed by the assertions of CLU [Liskov 1981] which, however, are not executable. Another CLU-based book by Liskov and Guttag [Liskov 1986], one of the few programming methodology texts to discuss in depth how to build reliable software, promotes the "defensive programming" approach of which the present chapter has developed a critique.

The notion of Design by Contract presented in this chapter and developed in the rest of this book comes from [M 1987a] and was further developed in [M 1988], [M 1989c], [M 1992b] and [M 1992c]. [M 1994a] discusses the tolerant and demanding approaches to precondition design, with particular emphasis on their application to the design of reusable libraries, and introduces the "tough love" policy. Further developments of the ideas have been contributed by James McKim in [McKim 1992a] (which led to some of the initial ideas for IFL), [McKim 1995], [McKim 1996], [McKim 1996a]; see also [Henderson-Sellers 1994a] which examines the viewpoint of the supplier.

EXERCISES

E11.1 Complex numbers

Write the abstract data type specification for a class **COMPLEX** describing the notion of complex number with arithmetic operations. Assume perfect arithmetic.

E11.2 A class and its ADT

Examine all the preconditions and axioms of the **STACK** abstract data type introduced in an earlier chapter and study whether and how each is reflected in class **STACK4**.

The ADT specification appeared on page 139. **STACK4** is on page 365 and includes **STACK4**, page 350

E11.3 Complete assertions for stacks

Show that by introducing a secret function *body* which returns the body of a stack, it is possible to make the assertions in a *STACK* class reflect the full corresponding abstract data type specification. Discuss the theoretical and practical value of this technique.

E11.4 Exporting the size

Page 350. Why is *capacity* exported for the bounded implementation of stacks, class *STACK2*?

E11.5 An implementation invariant

“A tolerant module”, page 360. Write the implementation invariant for class *STACK3*.

E11.6 Assertions and exports

Page 402. The discussion of using functions in assertions introduced a function *correct_index* for the precondition of *item* and *put*. If you add this function to class *ARRAY*, what export status must you give it?

E11.7 Finding the bugs

Page 382. Show that each of the four attempts at binary search algorithms advertised as “wrong” is indeed incorrect. (Hint: unlike proving an algorithm correct, which requires showing that it will work for all possible cases, proving it incorrect only requires that you find one case in which the algorithm will produce a wrong result, fail to terminate, or execute an illegal operation such as an out-of-bounds array access or other precondition violation.)

E11.8 Invariant violations

The discussion in this chapter has shown that a precondition violation indicates an error in the client, and a postcondition violation indicates an error in the supplier. Explain why an invariant violation also reflects a supplier error.

E11.9 Random number generators

Write a class implementing pseudo-random number generation, based on a sequence $n_i = f(n_{i-1})$ where f is a given function and the seed n_0 will be provided by clients of the class. Functions should have no side effects. (Assume f is known; you can find such functions in textbooks such as [Knuth 1981], and in numerical libraries.)

E11.10 A queue module

Write a class implementing queues (first-in, first-out policy), with appropriate assertions, in the style of the *STACK* classes of this chapter.

E11.11 A set module

Write a class implementing sets of elements of an arbitrary type, with the standard set operations (membership test, addition of a new element, union, intersection etc.). Be sure to include the proper assertions. Any correct implementation, such as linked lists or arrays, is acceptable.

POSTSCRIPT: THE ARIANE 5 CRASH

As this book was being printed, the European Space Agency released the report of the international investigation into the test flight of the Ariane 5 launcher, which crashed on June 4, 1996, 40 seconds after lift-off, at a reported cost of 500 million dollars (uninsured).

The cause of the crash: a failure of the on-board computer systems. The cause of that failure: a conversion from a 64-bit floating-point number (the mission’s “horizontal bias”) to a 16-bit signed integer produced an exception because the number was not representable with 16 bits. Although some other possible exceptions were monitored (using the Ada mechanisms described in the next chapter) prior analysis had shown that this particular one could not occur; so it was decided not to encumber the code with an extra exception handler.

The real cause: insufficient specification. The analysis that the value would always fit in 16 bits was in fact correct — but for the Ariane 4 flight trajectory! The code was reused for Ariane 5, and the assumption, although stated in an obscure part of some technical document, was simply forgotten. It did not apply any more to Ariane 5.

With the Design by Contract approach, it would have been stated in a precondition:

require

horizontal_bias <= *Maximum_horizontal_bias*

naturally prompting the quality assurance team to check all uses of the routine and to detect that some could violate the assertion. Although we will never know, it seems almost certain that the mistake would have been caught, probably through static analysis, and at worst, during testing thanks to the assertion monitoring mechanisms described in this chapter.

The lesson is clear: **reuse without contracts is folly**. The “abstracted modules” that we have defined as our units of reuse must be equipped with clear specifications of their operating conditions — preconditions, postconditions, invariants; and these specifications must be *in the modules themselves*, not in external documents. The principles that we have learned, particularly Design by Contract and Self-Documentation, are a required condition of any successful reusability policy. Even if your mistakes would cost less than half a billion dollars, remember this rule as you go after the great potential benefits of reuse: to be reusable, a module must be specified; and the programming language must support assertion mechanisms that will put the specification in the software itself.

For a more detailed discussion see [M 1997d].

When the contract is broken: exception handling

Like it or not, it is no use pretending: in spite of all static precautions, some unexpected and undesired event will sooner or later occur while one of your system's is executing. This is known as an exception and you must be prepared to deal with it.

12.1 BASIC CONCEPTS OF EXCEPTION HANDLING

The literature on exception handling is often not very precise about what really constitutes an exception. One of the consequences is that the exception mechanisms present in such programming languages as PL/I and Ada are often misused: instead of being reserved for truly abnormal cases, they end up serving as inter-routine *goto* instructions, violating the principle of Modular Protection.

Fortunately, the Design by Contract theory introduced in the preceding chapter provides a good framework for defining precisely the concepts involved.

Failures

Informally, an exception is an abnormal event that disrupts the execution of a system. To obtain a more rigorous definition, it is useful to concentrate first on a more elementary concept, failure, which follows directly from the contract idea.

A routine is not just some arbitrary sequence of instructions but the implementation of a certain specification — the routine's contract. Any call must terminate in a state that satisfies the precondition and the class invariant. There is also an implicit clause in the contract: that the routine must not have caused an abnormal operating system signal, resulting for example from memory exhaustion or arithmetic overflow and interrupting the normal flow of control in the system's execution.

It *must* refrain from causing such events, but of course not everything in life is what it must be, and we may expect that once in a while a routine call will be unable to satisfy its contract — triggering an abnormal signal, producing a final state that violates the postcondition or the invariant, or calling another routine in a state that does not satisfy that routine's precondition (assuming run-time assertion monitoring in the last two cases).

Such a case will be called a failure.

Definitions: success, failure

A routine call succeeds if it terminates its execution in a state satisfying the routine's contract. It fails if it does not succeed.

The discussion will use the phrase "routine failure", or just "failure", as an abbreviation for "failure of a routine call". Of course what succeeds or fails is not a routine (an element of the software text) but one particular call to that routine at run time.

Exceptions

From the notion of failure we can derive a precise definition of exceptions. A routine fails because of some specific event (arithmetic overflow, assertion violation...) that interrupts its execution. Such an event is an exception.

Definition: exception

An exception is a run-time event that may cause a routine call to fail.

Often an exception *will* cause failure of the routine. But you can prevent this from occurring by writing the routine so that it will catch the exception and try to restore a state from which the computation will proceed. This is the reason why failure and exception are different concepts: every failure results from an exception, but not every exception results in failure.

The study of software anomalies in the previous chapter introduced the terms *fault* (for a harmful execution event), *defect* (for an inadequacy of system, which may cause faults) and *error* (for a mistake in the thinking process, which may lead to defects). A failure is a fault; an exception is often a fault too, but not if its possible occurrence has been anticipated so that the software can recover from the exception.

See "Errors, defects and other creeping creatures", page 348.

Sources of exceptions

The software development framework introduced so far opens the possibility of specific categories of exception, listed at the top of the facing page.

Case E1 reflects one of the basic requirements of using references: a call *a.f* is only meaningful if *a* is attached to an object, that is to say non-void. This was discussed in the presentation of the dynamic model.

"Valid references and calls", page 246.

Case E2 also has to do with void values. Remember that "attachment" covers assignment and argument passing, which have the same semantics. We saw in the discussion of attachment that it is possible to attach a reference to an expanded target, the result being to copy the corresponding object. This assumes that the object exists; if the source is void, the attachment will trigger an exception.

"Hybrid attachments", page 263.

Definition: exception cases

An exception may occur during the execution of a routine *r* as a result of any of the following situations:

- E1 • Attempting a qualified feature call *a/* and finding that *a* is void.
- E2 • Attempting to attach a void value to an expanded target.
- E3 • Executing an operation that produces an abnormal condition detected by the hardware or the operating system.
- E4 • Calling a routine that fails.
- E5 • Finding that the precondition of *r* does not hold on entry.
- E6 • Finding that the postcondition of *r* does not hold on exit.
- E7 • Finding that the class invariant does not hold on entry or exit.
- E8 • Finding that the invariant of a loop does not hold after the **from** clause or after an iteration of the loop body.
- E9 • Finding that an iteration of a loop's body does not decrease the variant.
- E10 • Executing a **check** instruction and finding that its assertion does not hold.
- E11 • Executing an instruction meant explicitly to trigger an exception.

Case E3 follows from signals that the operating system sends to an application when it detects an abnormal event, such as a fault in an arithmetic operation (underflow, overflow) or an attempt to allocate memory when none is available.

Case E4 arises when a routine fails, as a result of an exception that happened during its own execution and from which it was not able to recover. This will be seen in more detail below, but be sure to note the rule that results from case E4:

Failures and exceptions

A failure of a routine causes an exception in its caller.

Cases E5 to E10 can only occur if run-time assertion monitoring has been enabled at the proper level: at least **assertion (require)** for E5, **assertion (loop)** for E8 and E9 etc.

Case E11 assumes that the software may include calls to a procedure **raise** whose sole goal is to raise an exception. Such a procedure will be introduced later.

Causes of failure

Along with the list of possible exception cases, it is useful for the record to define when a failure (itself the source of an exception in the caller, as per case E4) can occur:

See "Monitoring assertions at run time", page 393.

Definition: failure cases

A routine call will fail if and only if an exception occurs during its execution and the routine does not recover from the exception.

We have yet to see what it means for a routine to "recover" from an exception.

The definitions of failure and exception are mutually recursive: a failure arises from an exception, and one of the principal sources of exceptions is a calling routine (E4) is the failure of a called routine.

12.2 HANDLING EXCEPTIONS

We now have a definition of what may happen — exceptions — and of what we would prefer not to happen as a result — failure. Let us equip ourselves with ways to deal with exceptions so as to avoid failure. What can a routine do when its execution is suddenly interrupted by an unwelcome diversion?

As so often in this presentation, we can get help towards an answer by looking at examples of how *not* to do things. Here the C mechanism (coming from Unix) and an Ada textbook will oblige.

How not to do it — a C-Unix example

The first counter-example mechanism (most notably present on Unix, although it has been made available on other platforms running C) is a procedure called **signal** which you can call under the form

signal (signal_code, your_routine)

with the effect of planting a reference to *your_routine* into the software, as the routine that should be called whenever a signal of code *signal_code* occurs. A signal code is one of a number of possible integers such as **SIGILL** (illegal instruction) and **SIGFPE** (floating-point exception). You may include as many calls to **signal** as you like, so as to associate different routines with different signals.

Then assume some instruction executed after the call to **signal** triggers a signal of code *signal_code*. Were it not for the **signal** call, this event would immediately terminate the execution in an abnormal state. Instead it will cause a call to *your_routine*, which presumably performs some corrective action, and then will ... resume the execution exactly at the point where the exception occurred. This is dangerous, as you have no guarantee that the cause of the trouble has been addressed at all: if the computation was interrupted by a signal it was probably impossible to complete it starting from its initial state.

What you will need in most cases is a way to correct the situation and then **restart** the routine in a new, improved initial state. We will see a simple mechanism that implements this scheme. Note that one can achieve it in C too, on most platforms, by combining the **signal** facility with two other library routines: **setjmp** to insert a marker into the execution record for possible later return, and **longjmp** to return to such a marker, even if several calls have been started since the **setjmp**. The **setjmp-longjmp** mechanism is,

however, delicate to use; it can be useful in the target code generated by a compiler — and can indeed serve, together with *signal*, to implement the high-level O-O exception mechanism introduced later in this chapter — but is not fit for direct consumption by human programmers.

How not to do it — an Ada example

Here is a routine taken from an Ada textbook:

```

From Sommerville and Morrison, "Software Development with Ada", Addison-Wesley, 1987. Later case, indentation, semicolon usage and the name of the floating-point type have been adapted to the conventions of the present book. Non-positive has been changed to Negative.

sqrt (n: REAL) return REAL is
begin
    if x < 0.0 then
        raise Negative
    else
        normal_square_root_computation
    end
exception
    when Negative =>
        put ("Negative argument")
        return
    when others => ...
end -- sqrt

```

This example was probably meant just as a syntactical illustration of the Ada mechanism, and was obviously written quickly (for example it fails to return a value in the exceptional case); so it would be unfair to criticize it as if it were an earnest example of good programming. But it provides a useful point of reference by clearly showing an undesirable way of handling exceptions. Given the intended uses of Ada — military and space systems — one can only hope that not too many actual Ada programs have taken this model verbatim.

The goal is to compute the real square root of a real number. But what if the number is negative? Ada has no assertions, so the routine performs a test and, if it finds *n* to be negative, raises an exception.

The Ada instruction **raise** *Exc* interrupts execution of the current routine, triggering an exception of code *Exc*. Once raised, an exception can be caught, through a routine's (or block's) **exception** clause. Such a clause, of the form

```

exception
    when code_a1, code_a2, ... => Instructions_a;
    when code_b1, ... => Instructions_b;
    ...

```

is able to handle any exception whose code is one of those listed in the **when** subclauses; it will execute *Instructions_a* for codes *code_a1*, *code_a2*, ... and so on for the others. One of the subclauses may, as in the example, start with **when others**, and will then handle any exception not explicitly named in the other subclauses. If an exception occurs but its code

is not listed (explicitly or through **when others**), the routine will pass it to its caller; if there is no caller, meaning that the failed routine is the main program, execution terminates abnormally.

In the example there is no need to go to the caller since the exception, just after being raised, is caught by the **exception** clause of the routine itself, which contains a subclause **when Negative => ...**

But what then do the corresponding instructions do? Here they are again:

```

put ("Negative argument")
return

```

In other words: print out a message — a delicate thought, considering what happens next, and then return to the caller. The caller will not be notified of the event, and will continue its execution as if nothing had happened. Thinking again of typical applications of Ada, we may just wish that artillery computations, which can indeed require square root computations, do not follow this scheme, as it might direct a few missiles to the wrong soldiers (some of whom may, however, have the consolation of seeing the error message shortly before the encounter).

This technique is probably worse than the C-Unix *signal* mechanism, which at least picks up the computation where it left. A **when** subclause that ends with **return** does not even continue the current routine (assuming there are more instructions to execute); it gives up and returns to the caller as if everything were fine, although everything is *not* fine. Managers — and, to continue with the military theme, officers — know this situation well: you have assigned a task to someone, and are told the task has been completed — but it has not. This leads to some of the worst disasters in human affairs, and in software affairs too.

This counter-example holds a lesson for Ada programmers: under almost no circumstances should a **when** subclause terminate its execution with a **return**. The qualification "almost" is here for completeness, to account for a special case, the *false alarm*, discussed below; but that case is very rare. Ending exception handling with a **return** means pretending to the caller that everything is right when it is not. This is dangerous and unacceptable. If you are unable to correct the problem and satisfy the Ada routine's contract, you should make the routine fail. Ada provides a simple mechanism to do this: in an **exception** clause you may execute a **raise** instruction written as just

```
raise
```

whose effect is to re-raise the original exception to the caller. This is the proper way of terminating an execution that is not able to fulfill its contract.

Ada Exception rule

The execution of any Ada exception handler should end by either executing a **raise** instruction or retrying the enclosing program unit.

Exception handling principles

These counter-examples help show the way to a disciplined use of exceptions. The following principle will serve as a basis for the discussion.

Disciplined Exception Handling principle

There are only two legitimate responses to an exception that occurs during the execution of a routine:

R1 •Retrying: attempt to change the conditions that led to the exception and to execute the routine again from the start.

R2 •Failure (also known as **organized panic**): clean up the environment, terminate the call and report failure to the caller.

In addition, exceptions resulting from some operating system signals (case **E3** of the classification of exceptions) may in rare cases justify a **false alarm** response: determine that the exception is harmless and pick up the routine's execution where it started.

The classification of exception cases, including **E3**, is on page 413.

Let us do away first with the false alarm case, which corresponds to the basic C-Unix mechanism as we have seen it. Here is an example. Some window systems will cause an exception if the user of an interactive system resizes a window while some process is executing in it. Assume that such a process does not perform any window output; then the exception was harmless. But even in such case there are usually better ways, such as disabling the signals altogether, so that no exception will occur. This is how we will deal with false alarms in the mechanism of the next sections.

False alarms are only possible for operating system signals — in fact, only for signals of the more benign kind, since you cannot ignore an arithmetic overflow or an inability to allocate requested memory. Exceptions of all the other categories indicate trouble that cannot be ignored. It would be absurd, for example, to proceed with a routine after finding that its precondition does not hold.

So much for false alarms (unfortunately, since they are the easiest case to handle). For the rest of this discussion we concentrate on true exceptions, those which we cannot just turn off like an oversensitive car alarm.

Retrying is the most hopeful strategy: we have lost a battle, but we have not lost the war. Even though our initial plan for meeting our contract has been disrupted, we still think that we can satisfy our client by trying another track. If we succeed, the client will be entirely unaffected by the exception; after one or more new attempts following the initial failed one, we will return normally, having fulfilled the contract. (“Mission accomplished, Sir. The usual little troubles along the way, Sir. All fine by now, Sir.”)

What is the “other track” to be tried on the second attempt? It might be a different algorithm; or it might be the same algorithm, executed again after some changes have been brought to the state of the execution (attributes, local entities) in the hope of preventing the exception from occurring again. In some cases, it may even be the original routine tried again without any change whatsoever; this is applicable if the exception was due to some

external event — transient hardware malfunction, temporarily busy device or communication line — which we do not control although we expect it will go away.


With the other response, *failure*, we accept that we not only have lost the battle (the current attempt at executing the routine body) but cannot win the war (the attempt to terminate the call so as to satisfy the contract). So we give up, but we must first ensure two conditions, explaining the use of “organized panic” as a more vivid synonym for “failure”:

- Making sure (unlike what happened in the *sqrt* counter-example) that the caller gets an exception. This is the *panic* aspect: the routine has failed to live up to its contract.
- Restoring a consistent execution state — the *organized* aspect.

What is a “consistent” state? From our study of class correctness in the previous chapter we know the answer: a state that satisfies the invariant. We saw that in the course of its work a routine execution may temporarily violate the invariant, with the intention of restoring it before termination. But if an exception occurs in an intermediate state the invariant may be violated. The routine must restore it before returning control to its caller.

The call chain

To discuss the exception handling mechanism it will be useful to have a clear picture of the sequence of calls that may lead to an exception. This is the notion of call chain, already present in the explanation of the Ada mechanism.



The call chain

Let r_0 be the root creation procedure of a certain system (in Ada r_0 would be the main program). At any time during the execution, there is a *current routine*, the routine whose execution was started last; it was started by the execution of a certain routine; that routine was itself called by a routine; and so on. If we follow this called-to-caller chain all the way through we will end up at r_0 . The reverse chain (r_0 , the last routine r_1 that it called, the last routine r_2 that r_1 called, and so on down to the current routine) is the call chain.

If a routine produces an exception (as pictured at the bottom-right of the figure), it may be necessary to go up the chain until finding a routine that is equipped to handle the exception — or stop execution if we reach r_0 , not having found any applicable exception handler. This was the case in Ada when no routine in the call chain has an **exception** clause with a **when** clause that names the exception type or **others**.

12.3 AN EXCEPTION MECHANISM

From the preceding analysis follows the exception mechanism that fits best with the object-oriented approach and the ideas of Design by Contract.

The basic properties will follow from a simple language addition — two keywords — to the framework of the preceding chapters. A library class, *EXCEPTIONS*, will also be available for cases in which you need to fine-tune the mechanism.

Rescue and Retry

First, it must be possible to specify, in the text of a routine, how to deal with an exception that occurs during one of its calls. We need a new clause for that purpose; the most appropriate keyword is **rescue**, indicating that the clause describes how to try to recover from an undesirable run-time event. Because the **rescue** clause describes operations to be executed when the routine's behavior is outside of the standard case described by the precondition (**require**), body (**do**) and postcondition (**ensure**), it will appear, when present, after all these other clauses:

```
routine is
  require
  precondition
  local
  ... Local entity declarations ...
  do
  ... body
  ensure
  postcondition
  rescue
  rescue_clause
end
```

The *rescue_clause* is a sequence of instructions. Whenever an exception occurs during the execution of the normal *body*, this execution will stop and the *rescue_clause* will be executed instead. There is at most one **rescue** clause in a routine, but it can find out what the exception was (using techniques introduced later), so that you will be able to treat different kinds of exception differently if you wish to.

The other new construct is the **retry** instruction, written just **retry**. This instruction may only appear in a **rescue** clause. Its execution consists in re-starting the routine body from the beginning. The initializations are of course not repeated.

These constructs are the direct implementation of the Disciplined Exception Handling principle. The **retry** instruction provides the mechanism for retrying; a **rescue** clause that does not execute a **retry** leads to failure.

How to fail without really trying

The last observation is worth emphasizing:

Failure principle
Execution of a **rescue** clause to its end, not leading to a **retry** instruction, causes the current routine call to fail.

So if you have wondered how routines can fail in practice — causing case E4 of the exception classification — this is it. See page 413.

As a special case, consider a routine which does not have a **rescue** clause. In practice this will be the case with the vast majority of routines since the approach to exception handling developed here suggests equipping only a select few routines with such a clause. Ignoring possible local entity declarations, arguments, precondition and postcondition, the routine appears as

```
routine is
  do
  end
  body
end
Then if we consider — as a temporary convention — that the absence of a rescue
clause is the same thing as an empty rescue clause, that is to say
  routine is
  do
  body
  rescue
  — Nothing here (empty instruction list)
end
```

The Failure principle has an immediate consequence: if an exception occurs in a routine without **rescue** clause it will cause the routine to fail, triggering an exception in its caller.

Treating an absent **rescue** clause as if it were present but empty is a good enough approximation at this stage of the discussion; but we will need to refine this rule slightly when we start looking at the effect of exceptions on the class invariant. For the exact convention see “When there is no rescue clause”, page 430.

An exception history table

If a routine fails, either because it has no **rescue** clause at all or because its **rescue** clause executes to the end without a **retry**, it will interrupt the execution of its caller with a “Routine failed” (E4) exception. The caller is then faced with the same two possibilities: either it has a **rescue** clause that can execute a successful **retry** and get rid of the exception, or it will fail too, passing the exception one level up the call chain.

If in the end no routine in the call chain is able to recover from the exception, the execution as a whole will fail. In such a case the environment should print out a clear description of what happened, the exception history table. Here is an example:

An exception history table

Object	Class	Routine	Nature of exception	Effect
O4	Z.FUNCTION	split (from E.FUNCTION)	Feature <i>interpolate</i> : Called on void reference.	Retry
O3	INTERVAL	integrate	<i>interval_big_enough</i> : Precondition violated.	Fail
O2	EQUATION	solve (from GENERAL.EQUATION)	Routine failure	Fail
O2	EQUATION	filter	Routine failure	Retry
O2	MATH	new_matrix (from BASIC.MATH)	<i>enough_memory</i> : Check violated.	Fail
O1 (root)	INTERFACE	make	Routine failure	Fail

This is a record not only of the exceptions that directly led to the execution's failure but of all recent exceptions, up to a limit of 100 by default, including those from which the execution was able to recover through a *retry*. From top to bottom the order is the reverse of the order in which calls were started; the creation procedure is on the last line.


The **Routine** column identifies, for each exception, the routine whose call was interrupted by the exception. The **Object** column identifies the target of that call; here the objects have names such as O1, but in a real trace they will have internal identifiers, useful to determine whether two objects are the same. The **Class** column gives the object's generating class.

The **Nature of exception** column indicates what happened. This is where, for a assertion violation as in the second entry from the top, the environment can take advantage of assertion labels, *interval_big_enough* in the example, to identify the precise clause that was violated.

The last column indicates how the exception was handled: *Retry* or *Fail*. The table consists of a sequence of sections separated by thick lines; each section except the last led to a *Retry*. Since a *Retry* enables the execution to restart normally, an arbitrary number of calls may have occurred between two calls separated by a thick line.

Ignoring any such intermediate calls — successful and as such uninteresting for the purposes of this discussion — here is the call and return chain corresponding to the above exception history table. To reconstruct the action you should follow the arrows counter-clockwise from the call to *make* at the top left.

A failed execution



12.4 EXCEPTION HANDLING EXAMPLES

We now have the basic mechanism. Let us see how to apply it to common situations.

Fragile input

Assume that in an interactive system you need to prompt your system's user to enter an integer. Assume further that the only procedure at your disposal to read the integer, *read_one_integer*, leaving its result in the attribute *last_integer_read*, is not robust: if provided with something else than integer input, it may fail, producing an exception. Of course you do not want your own system to fail in that case, but since you have no control over *read_one_integer* you must use it as it is and try to recover from the exception if it occurs. Here is a possible scheme:

```
get_integer is
- Get integer from user and make it available in last_integer_read.
- If input initially incorrect, ask again as many times as necessary.
do
  print ("Please enter an integer:")
  read_one_integer
rescue
  retry
end
```

This version of the routine illustrates the retry strategy: we just keep retrying.

An obvious criticism is that if a user keeps on entering incorrect input, the routine will forever keep asking for a value. This is not a very good solution. We might put an upper bound, say five, on the number of attempts. Here is the revised version:

```
Maximum_attempts: INTEGER is 5
-- Number of attempts before giving up getting an integer.

get_integer is
-- Attempt to read integer in at most Maximum_attempts attempts.
-- Set value of integer_was_read to record whether successful.
-- If successful, make integer available in last_integer_read.
local
    attempts: INTEGER
do
    attempts := Maximum_attempts then
        print ("Please enter an integer: ")
        read_one_integer
        integer_was_read := True
    else
        integer_was_read := False
        attempts := attempts + 1
    end
    rescue
    retry
end

This assumes that the enclosing class has a boolean attribute integer_was_read
which will record how the operation went. Callers should use the routine as follows to try
to read an integer and assign it to an integer entity n:
```

```
get_integer
if integer_was_read then
    n := last_integer_read
else
    "Deal with case in which it was impossible to obtain an integer"
end
```

Recovering from hardware or operating system exceptions

Among the events that trigger exceptions are signals sent by the operating system, some of which may have originated with the hardware. Examples include: arithmetic overflow and underflow; impossible I/O operations; "illegal instruction" attempts (which, with a good object-oriented language, will come not from the O-O software but from companion routines, written in lower-level languages, which may overwrite certain areas of memory); creation or clone operations that fail because no memory is available; user interrupts (a user hitting the "break" key or equivalent during execution).

Theoretically you may view such conditions as assertion violations. If *a + b* provokes overflow, it means that the call has not observed the implicit precondition on the *+* function for integer or real numbers, stating that the mathematical sum of the two

arguments should be representable on the computer. A similar implicit precondition on the allocation of a new object (creation or clone) is that enough memory is available; if a write fails, it is because the environment — files, devices, users — did not meet the applicability conditions. But in such cases it is impractical or impossible to express the assertions, let alone check them: the only solution is to attempt the operation and, if the hardware or operating system signals an abnormal condition, to treat it as an exception.

Consider the problem of writing a function *quasi_inverse* which for any real number *x* must return either its inverse $\frac{1}{x}$ or, if that is impossible to compute because *x* is too small, the value 0. This type of problem is essentially impossible to solve without an exception mechanism: the only practical way to know whether *x* has a representable inverse is to attempt the division $\frac{1}{x}$; but if this provokes overflow and you cannot handle exceptions, the program will crash and it will be too late to return 0 as a result.

On some platforms it may be possible to write a function *invertible* such that *invertible(x)* is true if and only if the inverse of *x* can be computed. You can then use *invertible* to write *quasi_inverse*. But this is usually not a practical solution since such a function will not be portable across platforms, and in time-sensitive numerical computations will cause a serious performance overhead, a call to *invertible* being at least as expensive as the inversion itself.

With the **rescue-retry** mechanism you can easily solve the problem, at least on hardware that triggers a signal for arithmetic underflow:

```
quasi_inverse (x: REAL): REAL is
-- 1/x if possible, otherwise 0
local
    division_tried: BOOLEAN
do
    if not division_tried then
        Result := 1/x
    end
    rescue
        division_tried := True
    retry
end
```

Retrying for software fault tolerance

The initialization rules set *division_tried* to false at the start of each call. The body does not need any **else** clause because these rules also initialize *Result* to 0.

Assume you have written a text editor and (shame on you) you are not quite sure it is entirely bug-free, but you already want to get some initial user feedback. Your guinea pigs are willing to tolerate a system with some remaining errors; they might accept for example that once in a while it will be unable to carry out a command that they have requested; but

they will not use it to enter serious texts (which is what you want them to do, to test your editor under realistic conditions) if they fear that a failure may result in a catastrophe, such as brutal exit and loss of the last half-hour's work. With the `Retry` mechanism you can provide a defense against such behavior.

Assume that the editor, as will usually be the case, contains a basic command execution loop of the form

```
from ... until exit loop
    execute_one_command
end
```

where the body of routine `execute_one_command` is of the form

```
"Decode user request"
```

```
"Execute appropriate command in response to request"
```

The "Execute..." instruction chooses among a set of available routines (for example delete a line, change a word etc.) We will see in a later chapter how the techniques of inheritance and dynamic binding yield simple, elegant structures for such multi-way decisions.

The assumption is that the different routines are not entirely safe: some of them may fail at unpredictable times. You can provide a primitive but effective protection against such an event by writing the routine as

```
execute_one_command is
    -- Get a request from the user and, if possible,
    -- execute the corresponding command.
do
    "Decode user request"
    "Execute appropriate command in response to request"
rescue
    message ("Sorry, this command failed")
    message ("Please try another command")
    message ("Please report this failure to the author")
    "Instructions to patch up the state of the editor"
retry
end
```

This scheme assumes in practice that the types of supported `user request` include "save current state of my work" and "quit", both of which had better work correctly. A user who sees the message *Sorry, this command failed* will most likely want to save the results of the current session and get out as quickly as possible.

Some of the routines implementing individual operations may have their own `rescue` clauses, leading to failure (so that the above `rescue` clause of `execute_one_command` takes over) but only after printing a more informative, command-specific message.

N-version programming

Another example of retrying for software fault tolerance is an implementation of the "N-version programming" approach to improving software reliability.

N-version programming was inspired by redundancy techniques that have proved their usefulness in hardware. In mission-critical setups it is frequent to encounter redundant hardware configurations, where several devices — for example computers — perform an identical function, and an arbitrating mechanism compares the results, deciding for the majority in case of discrepancy. This approach guards against single-component failures and is common in aerospace applications. (In a famous incident, an early space shuttle launch had to be delayed because of a bug in the software for the *arbitrating* computer itself.) N-version programming transposes this approach to software by suggesting that for a mission-critical development two or more teams, working in environments as distinct as possible, should produce alternative systems, in the hope that errors, if any, will be different.

This is a controversial idea: one may argue that the money would be better spent in improving the correctness and robustness of a single version than in financing two or more imperfect implementations. Let us, however, ignore these objections and refrain from any judgment on the idea itself, but see how the `retry` mechanism would support the idea of using several implementations where one takes over if the others fail:

```
do_task is
    -- Solve a problem by applying one of several possible implementations.
    require
    ...
    local
        attempts: INTEGER
    do
        if attempts = 0 then
            implementation_1
        elsif attempts = 1 then
            implementation_2
        end
        ensure
        ...
    rescue
        attempts := attempts + 1
        if attempts < 2 then
            "Perhaps some instructions to reset to stable state"
            retry
        end
    end
end
```

The generalization to more than two alternative implementations is immediate.

This example is typical of the use of `retry`. The `retry` clause *never* attempts to reach the original goal using a substitute implementation; reaching this goal, as expressed by the `postcondition` if there is one, is the privilege of the normal body. Note that after two attempts (or `n` in the general case) the routine simply executes its `retry` clause to the end, and so fails.

Let us look more closely at what happens when an exception is triggered during the execution of `r`. The normal execution (the body) stops; the rescue clause is executed instead. Then two cases may occur:

- The rescue clause may execute a **retry**, usually after some other instructions. In this case, execution of the routine will start anew. This new attempt may succeed; in this case, execution of the routine will terminate normally and return to its client. The call is a success; the contract has been fulfilled. Execution of the client is not affected, except of course that the call may have taken longer than normal. If, however, the **retry** attempt again fails, the rescue clause will execute again. This time, however, the **retry** attempt causes an exception; the process of executing the rescue clause will start anew.

- If the rescue clause does not execute a **retry**, it will continue to its end. (This happens in the last example when *attempts > 2*.) In this case the routine fails: it returns control to its caller, signaling an exception. Because the caller gets an exception, the same rule determines how its own execution continues.

This mechanism strictly adheres to the Disciplined Exception Handling principle: either a routine succeeds, that is to say its body executes to the end and satisfies the postcondition, or it fails. When interrupted by an exception, you may either report failure or try your normal body again; in no way can you exit through the rescue clause and pretend to your caller that you succeeded.


12.5 THE TASK OF A RESCUE CLAUSE

The last comments get us started towards a better understanding of the exception mechanism by suggesting the theoretical role of rescue clauses. Some formal reasoning will help us obtain the complete picture.

The correctness of a rescue clause

See “WHEN IS A CLASS CORRECT?” *RECT’97*, 11.9, page 370.

The formal definition of class correctness stated two requirements on the features of a class. One (C1) requires creation procedures to start things off properly. The other (C2), more directly relevant for the present discussion, states that to satisfy its contract, every routine, started with its precondition and the class invariant both satisfied, must preserve the invariant and ensure its postcondition. This was illustrated by the diagram depicting the typical object lifecycle:



The life of an object

The formal rule read:

C2 • For every exported routine r and any set of valid arguments x_r :

$\{pre_r(x_r)$ and $INV\}$ Body $\{post_r(x_r)$ and $INV\}$

This rule appeared on page 371.

where *pre_r* is the precondition, *INV* the class invariant, *Body_r* the body of the routine, and *post_r* the postcondition. To keep things simple let us ignore the arguments X_r .

Let *Rescue*_{*i*} be the rescue clause of a routine, ignoring any branch that leads to a *retry* — that is to say keeping only those branches that will result in failure if executed. Rule C2 is a specification of the body *Body_i* of the routine, in terms of what initial states it assumes and what final states it can guarantee. Can we obtain a similar specification for *Rescue_i*? It should be of the form

$$\{\cdot\} \text{ Rescue } \{\cdot\}$$

with the question marks replaced by actual assertions. (Here it is useful to try answering the question for yourself before reading on: how would you fill in the question marks?)

Consider first the input assertion — the question mark on the left of *Rescue*_. Any non-trivial (that we write there would be wrong!) Remember the discussion of attractive job offers: for whoever implements the task A in $(P) \vdash A \{Q\}$, the stronger the precondition P , the easier the job, since a precondition restricts the set of input cases that you must handle. Any precondition for *Rescue*_* would make the job easier by restricting the set of states in which *Rescue*_* may be called to action. But we may not assume any such restriction since exceptions, by their very nature, may happen at any time. If we knew when an exception will happen, it probably would not be an exception any more. Think of hardware failures: we have no clue as to when a computer can start to malfunction. Nor do we know, in an interactive system, when a user will feel like hitting the “break” key.

So the only *P* assertion that we can afford here (to replace the question mark on the left) is the one that asserts nothing at all: *True*, the assertion that all states satisfy.

For a lazy *Rescue*, implementor — again in reference to the discussion of job offers in the previous chapter — this is bad news; in fact the precondition *True* is always the worst possible news for a supplier, the case in which “the customer is always right”!

What about the output assertion (the *Q*)? As discussed earlier, a rescue clause that leads to a failure must, before returning control to its caller with an exception, restore a stable state. This means reestablishing the invariant.

Hence the rule that we seek, with no more question marks:

C3 •

{True} Rescue, {INV}

Correctness rule for failure-inducing rescue clauses

Similar reasoning yields the corresponding rule for any branch *Retry*, of the rescue clause leading to a *retry* instruction:

C4 •

{True} Retry, {INV and pre_r}

Correctness rule for retry-inducing rescue clauses

A clear separation of roles

It is interesting to contrast the formal roles of the body and the rescue clause:

- C2 • {pre_r and INV} Body, {post_r (x_r) INV}
- C3 • {True} Rescue, {INV}

The input assertion is stronger for *Body_r*; whereas the rescue clause is not permitted to assume anything at all, the routine's body may assume the precondition and the invariant. This makes its job easier.

The output assertion, however, is also stronger for *Body_r*; whereas the rescue clause is only required to restore the invariant, the normal execution must also ensure the postcondition — the official job of the routine. This makes its job harder.

These rules reflect the separation of roles between the body (the *do* clause) and the rescue clause. The task of the body is to ensure the routine's contract; not directly to handle exceptions. The task of the rescue clause is to handle exceptions, returning control to the body or (in the failure case) to the caller, not to ensure the contract.

As an analogy — part of this book's constant effort to provide readers not just with theoretically attractive concepts but also with practical skills that they can apply to the pursuit of their careers — consider the difficulty of choosing between two noble professions: *cook* and *firefighter*. Each has its grandeur, but each has its servitudes. A gratifying quality of the cook's job is that he may assume, when he shows up at work in the morning, that the restaurant is not burning (satisfies the invariant); presumably his contract does not specify any cooking obligation under burning circumstances. But with a non-burning initial state the cook must prepare meals (ensure the postcondition); it is also a component of his contract, although perhaps an implicit one, that throughout this endeavor he should maintain the invariant, if he can, by not setting the restaurant on fire.

The firefighter, for his part, may assume nothing as to the state in which he finds the restaurant when he is called for help at any time of day or night. There is not even any guarantee that the restaurant is indeed burning — no precondition of the form *is burning*, or of any other form save for *True* — since any call may be a false alarm. In some cases, of course, the restaurant will be burning. But then a firefighter's only duty is to return it to a non-burning state; his job description does not require that he also serve a meal to the assembly of patiently waiting customers.

When there is no rescue clause

Having formalized the role of rescue clauses we can take a second look at what happens when an exception occurs in a routine that has no such clause. The rule introduced earlier — with a warning that it would have to be revised — stated that an absent rescue clause was equivalent to a present but empty one (*rescue end*). In light of our formal rules, however, this is not always appropriate. C3 requires that

{True} Rescue, {INV}

If *Rescue_r* is an empty instruction and the invariant *INV* is anything other than *True*, this will not hold.

Hence the exact rule. The class *ANY* — mother of all classes — includes a procedure

```
default_rescue is
- Handle exception if no Rescue clause.
- (Default: do nothing)
end
```

We will study *ANY*, whose features are present in all classes, in "THE GLOBAL INHERITANCE STRUCTURE", 162, page 580.

A routine that does not have a Rescue clause is considered to have one that, rather than being empty as first suggested, has the form

```
rescue
default_rescue
```

Every class can redefine *default_rescue* (using the feature redefinition mechanism studied as part of inheritance in a later chapter) to perform some specific action, instead of the default empty effect defined in *GENERAL*.

Rule C3 indicates the constraint on any such action: starting in any state, it must restore the class invariant *INV*. Now you will certainly remember that producing a state that satisfies the invariant was also the role of the **creation procedures** of a class, as expressed by the rule labeled C1. In many cases, you will be able to write the redefinition of *default_rescue* so that it relies on a creation procedure.

See "The role of creation procedures", page 372, Rule C1 was on page 371.

12.6 ADVANCED EXCEPTION HANDLING

The extremely simple mechanism developed so far handles most of the needs for exception handling. But certain applications may require a bit of fine-tuning:

- You may need to find out the nature of the latest exception, so as to handle different exceptions differently.
- You may want to specify that certain signals should not trigger an exception.
- You may decide to trigger an exception yourself.

We could extend the language mechanism accordingly, but this does not seem the right approach, for at least three reasons: the facilities are needed only occasionally, so that we would be needlessly burdening the language; some of them (in particular anything that has to do with signals) may be platform-dependent, whereas a language definition should be portable; and when you select a set of these facilities it is hard to be sure that you will not at some later time think of other useful ones, which would then force a new language modification — not a pleasant prospect.

For such a situation we should turn not to the language but to the supporting library. We introduce a library class *EXCEPTIONS*, which provides the necessary fine-tuning capabilities. Classes that need these capabilities will inherit *EXCEPTIONS*, using the inheritance mechanism detailed in later chapters. (Some developers may prefer to use the client relation rather than inheritance.)

Exception queries

Class *EXCEPTIONS* provides a number of queries for obtaining some information about the last exception. You can find out the integer code of that exception:

```
exception: INTEGER
-- Code of last exception that occurred

original_exception: INTEGER
-- Original code of last exception that triggered current exception
```

The difference between *exception* and *original_exception* is significant in the case of an “organized panic” response: if a routine gets an exception of code *oc* (indicating for example an arithmetic overflow) but has no rescue clause, its caller will get an exception whose own code, given by the value of *exception*, indicates “failure of a called routine”. It may be useful at that stage, or higher up in the call chain, to know what the original cause was. This is the role of *original_exception*.

The exception codes are integers. Values for the predefined exceptions are given by integer constants provided by *EXCEPTIONS* (which inherits them from another class *EXCEPTION_CONSTANTS*). Here are some examples:

```
Check_instruction: INTEGER is 7
-- Exception code for violated check

Class_invariant: INTEGER is ...
-- Exception code for violated class invariant

Incorrect_inspect_value: INTEGER is ...
-- Exception code for inspect value which is not one
-- of the inspect constants, if there is no Else_part

Loop_invariant: INTEGER is ...
-- Exception code for violated loop invariant

Loop_variant: INTEGER is ...
-- Exception code for non-decreased loop variant

No_more_memory: INTEGER is ...
-- Exception code for failed memory allocation

Postcondition: INTEGER is ...
-- Exception code for violated postcondition

Precondition: INTEGER is ...
-- Exception code for violated precondition

Routine_failure: INTEGER is ...
-- Exception code for failed routine

Void_assigned_to_expanded: INTEGER is ...

Since the integer values themselves are irrelevant, only the first one has been shown.

A few other self-explanatory queries provide further information if needed:

meaning (except: INTEGER)
-- A message in English describing the nature of exceptions
-- of code except

is_assertion_violation: BOOLEAN
-- Is last exception originally due to a violated assertion
-- or non-decreasing variant?

ensure
Result = (exception = Precondition) or (exception = Postcondition) or
(exception = Class_invariant) or
(exception = Loop_invariant) or (exception = Loop_variant)

is_system_exception: BOOLEAN
-- Is last exception originally due to external event (operating system error)?

is_signal: BOOLEAN
-- Is last exception originally due to an operating system signal?

tag_name: STRING
-- Tag of last violated assertion clause

original_tag_name: STRING
-- Assertion tag for original form of last assertion violation.
```

```
recipient_name: STRING
    - Name of routine whose execution was interrupted by last exception
class_name: STRING
    - Name of class including recipient of last exception
original_recipient_name: STRING
    - Name of routine whose execution was interrupted by
    - original form of last exception
original_class_name: STRING
    - Name of class including recipient of original form of last exception
```

With these features, a rescue clause can handle different kinds of exception in different ways. For example you can write it, in a class inheriting from *EXCEPTIONS*, as

```
rescue
    if is_assertion_violation then
        "Process assertion violation case"
    else if is_signal then
        "Process signal case"
    else
        ...
    end
or, with an even finer grain of control, as
```

```
rescue
    if exception = Incorrect_inspect_value then
        "Process assertion violation case"
    else if exception = Routine_Failure then
        "Process signal case"
    else
        ...
    end
```

quasi_inverse was on page 424.

Using class *EXCEPTIONS*, we can modify the *quasi_inverse* example so that it will only attempt the *retry*. If the exception was overflow. Other exceptions, such as one generated when the interactive user presses the Break key, will not cause the *retry*. The instruction in the rescue clause becomes:

```
if exception = Numerical_error then
    division_tried := True; retry
end
```

Since there is no *else* clause, exceptions other than *Numerical_error* will result in failure, the correct consequence since the routine has no provision for recovery in such cases. When writing a rescue clause specifically to process a certain kind of possible exception, you may use this style to avoid lumping other, unexpected kinds with it.

How fine a degree of control?

One may express reservations about going to the level of specific exception handling illustrated by the last two extracts. This chapter has developed a view of exceptions as undesired events; when one happens, the reaction of the software and its developer is "I don't want to be here! Get me out as soon as possible!". This seems incompatible with exerting fine control, depending on the exception's source, over what happens in a rescue clause.

For that reason, I tend in my own work to avoid using detailed case analysis on exception sources, and stick to exception clauses that try to fix things if they can, and then fail or *retry*.

This style is perhaps too austere, and some developers prefer a less restricted exception handling scheme that makes full use of the query mechanisms of class *EXCEPTIONS* while remaining disciplined. If you want to use such a scheme you will find in *EXCEPTIONS* all that you need. But do not lose sight of the following principle, a consequence of the discussion in the rest of this chapter:

Exception Simplicity principle

All processing done in a rescue clause should remain simple, and focused on the sole goal of bringing the recipient object back to a stable state, permitting a *retry* if possible.

Developer exceptions

All the exceptions studied so far resulted from events caused by agents external to the software (such as operating system signals) or from involuntary consequences of the software's action (as with assertion violations). It may be useful in some applications to cause an exception to happen on purpose.

Such an exception is called a developer exception and is characterized by an integer code (separate from the general exception code, which is the same for all developer exceptions) and an associated string name, which may be used for example in error messages. You can use the following features to raise a developer exception, and to analyze its properties in a rescue clause:

```
trigger (code: INTEGER; message: STRING)
    - Interrupt execution of current routine with exception
    - of code code and associated text message.

developer_exception_code: INTEGER
developer_exception_name: STRING
    - Name associated with last developer exception
```

```

is_developer_exception: BOOLEAN
    -- Was last exception originally due to a developer exception?

is_developer_of_name (name: STRING): BOOLEAN
    -- Is the last exception originally due to a developer
    -- exception of name name?

```

```

ensure
    Result := is_developer_exception and then
        equal (name, developer_exception_name)

```

It is sometimes useful to associate with a developer exception a context — any object structure that may be useful to the software element handling the exception:

```

set_developer_exception_context (c: ANY)
    -- Define c as the context associated with subsequent developer
    -- exceptions (as caused by calls to trigger).

require
    context_exists: c /= Void
    developer_exception_context: ANY
    -- Context set by last call to set_developer_exception_context
    -- void if no such call.

```

These facilities enable a style of development that heavily relies on some software elements triggering exceptions that others will process. In one compiler that I have seen, the developers took advantage of this mechanism, in the parsing algorithm, to stick to a relatively linear control structure, recognizing the various elements of the input text one after the other. Such sequential treatment is only possible if the elements parsed are the expected ones; any syntactical error disrupts the process. Rather than complicating the control structure by adding possibly nested **if ... then ... else** constructs, the developers chose to raise a developer exception whenever encountering an error, then dealt with it separately in the calling routine. As hinted earlier, this is not my favorite style, but there is nothing inherently wrong with it, so the developer exception mechanisms are there for those who want them.

12.7 DISCUSSION

We have now completed the design of an exception mechanism for object-oriented software construction, compatible with the rest of the approach and resulting directly from the ideas of Design by Contract developed in the preceding chapter. Thanks in particular to the **retry** instructions the mechanism is more powerful than what you will find in many languages, at the same time it may appear stricter because of its emphasis on retaining the ability to reason precisely about the effect of each routine.

Let us explore a few alternative design ideas that could have been followed, and why they were not retained.

Disciplined exceptions

Exceptions, as they have been presented, are a technique to deal with erroneous conditions that may be arise at run time: assertion violations, hardware signals, attempts to access void references.

The approach we have explored is based on the contracting metaphor: under no circumstances should a routine pretend it has succeeded when in fact it has failed to achieve its purpose. A routine may only succeed (perhaps after experiencing some exceptions but recovering from them through one or more **retry**, unbeknownst to the client) or fail.

Exceptions in Ada, C++ or PL/I do not follow this model. Using the Ada model, and instruction

```

raise exc

```

cancels the routine that executed it and returns control to its caller, which may handle the exception **exc** in a special handler clause or, if it has no such handler, will itself return control to its caller. But there is no rule as to what a handler may do. Hence it is perfectly possible to ignore an exception, or to return an alternate result. This explains why some developers use this exception mechanism simply to deal with cases other than the easiest one for an algorithm. Such applications of exceptions really use **raise** as a **goto** instruction, and a fairly dangerous one since it crosses routine boundaries. In my opinion, they are abuses of the mechanism.

There have traditionally been two viewpoints on exceptions: many practicing programmers, knowing how essential it is to retain control at run time whenever an abnormal condition is detected (whether due to a programming error or to an unforeseeable hardware event, say numerical overflow or hardware failure), consider them an indispensable facility. Computing scientists preoccupied with correctness and systematic software construction have often for their part viewed exceptions with suspicion, as an unclean facility used to circumvent the standard rules on control structures. The mechanism developed above will, it is hoped, appeal to both sides.

Should exceptions be objects?

An object-oriented zealot (and who, having discovered and mastered the beauty of the approach, does not at times risk succumbing to zeal?) may criticize the mechanism presented in this chapter for not treating exceptions as first-class citizens of our software society. Why is an exception not an object?

One recent language, the object-oriented extension of Pascal for Borland's Delphi environment has indeed taken the attitude that exceptions should be treated as objects.

It is not clear that such a solution would bring any benefit. The reasoning is in part a preview of the more general discussion that will help us, in a later chapter, tackle the question “how do we find the objects and classes?” An object is an instance of an abstractly defined data type, characterized by features. An exception has some features, of

^{“No-command classes”, page 729, See exercise E12.2, page 438.}

course, which we saw in class *EXCEPTIONS*: its type, which was given by an integer code; whether it is a signal, an assertion violation, a developer exception; its associated message if it is a developer exception. But these features are *queries*: in most classes describing true “objects” there should also be *commands* changing the objects’ state. Although one might conceive of commands applicable to exception objects, for example to disarm an exception after it has occurred, this seems incompatible with reliability requirements. Exceptions are not under the control of the software system; they are triggered by events beyond its reach.

Making their properties accessible through the simple queries and commands of the class *EXCEPTIONS* seems enough to satisfy the needs of developers who want fine-grain access to the exception handling mechanism.

The methodological perspective

“DEALING WITH ABNORMAL CASES”, 23.6, page 797.
A final note and preview. Exception handling is not the only answer to the general problem of *robustness* — how to deal with special or undesired cases. We have gained a few methodological insights, but a more complete answer will appear in the chapter discussing the design of module interfaces, allowing us to understand the place of exception handling in the broader arsenal of robustness-enhancing techniques.

12.8 KEY CONCEPTS INTRODUCED IN THIS CHAPTER

- Exception handling is a mechanism for dealing with unexpected run-time conditions.
- A failure is the impossibility, for a routine execution, to fulfill the contract.
- A routine gets an exception as a result of the failure of a routine which it has called, of an assertion violation, of an abnormal condition signaled by the hardware or operating system.
- It is also possible for a software system to trigger a “developer exception” explicitly.
- A routine will deal with an exception by either *Retry* or *Organized Panic*. *Retry* reexecutes the body; *Organized Panic* causes a routine failure and sends an exception to the caller.
- The formal role of an exception handler not ending with a *retry* is to restore the invariant — not to ensure the routine’s contract, as that is the task of the body (the *do* clause). The formal role of a branch ending with *retry* is to restore the invariant and the precondition so that the routine body can try again to achieve its contract.
- The basic language mechanism for handling exceptions should remain simple, if only to encourage straightforward exception handling — organized panic or retrying. For applications that need finer control over exceptions, their properties and their processing, a library class called *EXCEPTIONS* is available: it provides a number of mechanisms for distinguishing between exception types, as well as for triggering developer-defined exceptions.

12.9 BIBLIOGRAPHICAL NOTES

[Liskov 1979] and [Christian 1985] offer other viewpoints on exceptions. Much of the work on software fault tolerance derives from the notion of “recovery block” [Randell 1975]: a recovery block for a task is used when the original algorithm for the task fails to succeed. This is different from rescue clauses which never by themselves attempt to achieve the original goal, although they may restart the execution after patching up the environment.

[Hoare 1981] contains a critique of the Ada exception mechanism.

The approach to exception handling developed in this chapter was first presented in [M 1988e] and [M 1988f].

EXERCISES

E12.1 Largest integer

Assume a machine that generates an exception when an integer addition overflows. Using exception handling, write a reasonably efficient function that will return the largest positive integer representable on the machine.

E12.2 Exception objects

Notwithstanding the skeptical comments expressed in the discussion section as to the usefulness of treating exceptions as objects, press the idea further and discuss what a class *EXCEPTION* would look like, assuming an instance of that class denotes an exception that has occurred during execution. (Do not confuse this class with *EXCEPTIONS*, the class, meant to be used through inheritance, which provides general exception properties.) Try in particular to include commands as well as queries.