Downloaded from rsta.royalsocietypublishing.org

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Mechanical Proofs about Computer Programs [and Discussion]

D. I. Good and B. A. Wichmann

Phil. Trans. R. Soc. Lond. A 1984 **312**, 389-409
doi: 10.1098/rsta.1984.0066

Email alerting service Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

To subscribe to *Phil. Trans. R. Soc. Lond. A* go to: **http://rsta.royalsocietypublishing.org/subscriptions**

This journal is © 1984 The Royal Society

Downloaded from rsta.royalsocietypublishing.org

Phil. Trans. R. Soc. Lond. A **312**, 389–409 (1984) [389]
Printed in Great Britain

Mechanical proofs about computer programs

By D. I. Good

2100 *Main Building, Institute for Computing Science, The University of Texas at Austin, Austin, Texas* 78712, *U.S.A.*

The Gypsy verification environment is a large computer program that supports the development of software systems and formal, mathematical proofs about their behaviour. The environment provides conventional development tools, such as a parser for the Gypsy language, an editor and a compiler. These are used to evolve a library of components that define both the software and precise specifications about its desired behaviour. The environment also has a verification condition generator that automatically transforms a software component and its specification into logical formulas that are sufficient to prove that the component always runs according to specification. Facilities for constructing formal, mechanical proofs of these formulas also are provided. Many of these proofs are completed automatically without human intervention. The capabilities of the Gypsy system and the results of its application are discussed.

1. Introduction

One of the major problems with the current practice of software engineering is an absence of predictability. There is no sound, scientific way of predicting accurately how a software system will behave when it runs. There are many compelling examples of important software systems that have behaved in unpredictable ways: a Space Shuttle fails to launch; an entire line of automobiles is recalled because of problems with the software that controls the braking system; unauthorized users get access to computer systems; sensitive information passes into the wrong hands, etc. (Neumann 1983*a*, *b*). Considering the wide variety of tasks that now are entrusted to computer systems, it is truly remarkable that it is not possible to predict accurately what they are going to do!

Within current software engineering practice, the only sound way to make a precise, accurate prediction about how a software system will behave is to build it and run it. There is no way to predict accurately how a system will behave before it can be run. So design flaws often are detected only after a large investment has been made to develop the system to a point where it can be run. The rebuilding that is caused by the late detection of these flaws contributes significantly to the high cost of software construction and maintenance. Even after the system can be run, the situation is only slightly better. A system that can be run can be tested on a set of trial cases. If the system is deterministic, a trial run on a specific test case provides a precise, accurate prediction about how the system will behave in *that one case*. If the system is re-run on exactly the same case, it will behave in exactly the same way. However, there is no way to predict, from the observed behaviour of a finite number of test cases, how the system will behave in any other case. If the system is non-deterministic (as many systems are), the system will not even necessarily repeat its observed behaviour on a test case. So in current software engineering practice, predicting that a software system will run according to specification is based almost entirely on subjective, human judgment rather than on objective, scientific fact.

In contrast to software engineering, mathematical logic provides a sound, objective way to

The Royal Society is collaborating with JSTOR to digitize, preserve, and extend access to *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences.* www.jstor.org

Downloaded from rsta.royalsocietypublishing.org

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

make accurate, precise predictions about the behaviour of mathematical operations. For example, if x and y are natural numbers, who among us would doubt the prediction that $x+y$ always gives exactly the same result as $y+x$? This prediction is accurate not just for *some* cases or even just for *most* cases; it is accurate for *every* pair of natural numbers, no matter what they are. The prediction is accurate because there is a *proof* that $x+y = y+x$ logically follows from accepted definitions of 'natural number', '=' and '+'.

The Gypsy verification environment is a large, interactive computer program that supports the construction of formal, mathematical proofs about the behaviour of software systems. These proofs make it possible to predict the behaviour of a software system with the same degree of precision and accuracy that is possible for mathematical operations. These proofs can be constructed *before* a software system can be run and, therefore, they can provide an objective, scientific basis for making predictions about system behaviour throughout the software life cycle. This makes it possible for the proofs actually to guide the construction of the system. In theory these proof methods make possible a new approach to software engineering that can produce systems whose predictability far exceeds that which can be attained with conventional methods.

In practice the use of this mathematical approach to software engineering requires very careful management of large amounts of detailed information. The Gypsy environment is an experimental system that has been developed to explore the viability of applying these methods in actual practice. The purposes of the environment are to amplify the ability of the human software engineer to manage these details and to reduce the probability of human error. The environment, therefore, contains tools for supporting the normal software development process as well as tools for constructing formal proofs.

2. A mathematical approach

The Gypsy verification environment is based on the Gypsy language (Good *et al.* 1978). Rather than being based on an extension of the hardware architecture of some particular computer, the Gypsy language is based on rigorous, mathematical foundations for specifying and implementing computer programs. The specification describes *what* effect is desired when the program runs, and the implementation defines *how* the effect is caused. The mathematical foundation provided by the Gypsy language makes it possible to construct rigorous proofs about both the specifications and the implementations of software systems. The language, which is modelled on Pascal (Jensen & Wirth 1974), also is designed so that the implementations of programs can be compiled and executed on a computer with a conventional von Neumann architecture.

The basic structure of a Gypsy software system is shown in figure 1. The purpose of a software system is to cause some effect on its external environment. The external environment of a Gypsy software system consists of data objects (and exception conditions). Every Gypsy data object has a name and a value. The implementation of a program causes an effect by changing the values of the data objects in its external environment (or by signalling a condition). To accomplish its effect, an implementation may create and use internal (local) data objects (and conditions). In figure 1, X and Y represent external objects, and U represents an internal object.

The specifications of a program define constraints on its implementation. In parallel with the structure of implementations, Gypsy provides a means of stating both internal and external specifications. The external specifications constrain the externally visible effects of an implementation. Internal specifications constrain its internal behaviour.

Downloaded from rsta.royalsocietypublishing.org

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

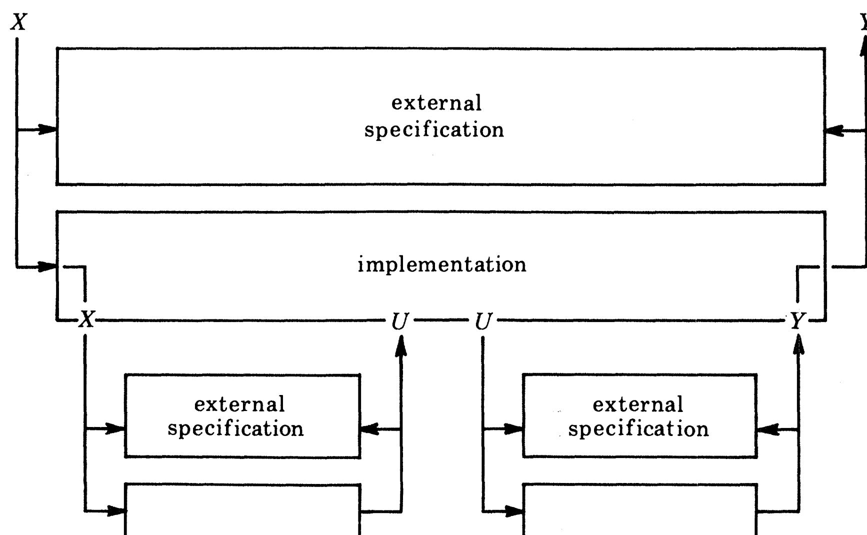

FIGURE 1. Gypsy software system structure.

The external specifications of a program consist of two parts: a (mandatory) environment specification and an (optional) operational specification. The environment specification describes *all* of the external data objects that are accessible to the procedure. The specification also states the type of each of these objects and whether it is a variable or a constant object. A program may change the value of a data object *only* if it is a variable object.

The type of an object specifies the kind of values it may have. The mathematical foundations of Gypsy begin with its types. The Gypsy types are all well known mathematical objects (integers, rational numbers, the Boolean values *true* and *false*, sets, sequences, mappings) or they can be easily derived from such objects (types character, record, array, buffer). For example, in Gypsy, type *integer* represents the full, unbounded set of mathematical objects. It is not restricted only to the integers that can be represented on a particular machine. For each of these pre-defined types, the Gypsy language also provides a set of primitive, pre-defined functions with known mathematical properties.

The operational specification for an implementation is a relation (a Boolean-valued function) that describes what effect is to be caused on the objects of the external environment. These relations are defined by ordinary functional composition from the Gypsy pre-defined functions.

The implementation of a Gypsy program is defined by a procedure. Running a Gypsy procedure is what actually causes an effect to be produced in its external environment. For implementation, the Gypsy language provides a set of pre-defined procedures (assign a value to an object, send a value to a buffer, remove an object from a sequence, etc.) that have precisely defined effects. It also provides a set of composition rules (*if...then...else...end*, *loop...end*, *cobegin...end*, etc.) for composing these pre-defined procedures into more complex ones. So the implementation of every Gypsy software system is some composition of the pre-defined procedures.

These composition rules are designed so that the effect that is caused by the composition can be deduced from the effects caused by its components. In particular, it is always possible to construct a set of formulas in the first-order predicate calculus that are sufficient (but not always necessary) to show that the effect caused by a procedure satisfies its specifications. These formulas are called *verification conditions*. They are the logical conditions that are sufficient to

Downloaded from rsta.royalsocietypublishing.org

verify that the implementation meets its specifications. By constructing them, the task of proving that an implementation always causes an effect that satisfies its specifications is reduced to a task of proving a set of formulas in the first-order predicate calculus. The methods for constructing the verification conditions are based on the pioneering work of Naur (1966), Floyd (1967), Dijkstra (1968), Hoare (1969), King (1969), Good (1970). Dijkstra (1976), Jones (1980), Gries (1981), Hoare (1982) provide more recent discussions of these basic ideas and their relation to software development.

One of the most important aspects of the Gypsy composition rules is illustrated in figure 1. Only the external specifications of the components are required to construct the verification conditions for the composition. Neither the internal specifications nor the implementation of the components are required. The proof of the composition is completely independent of the internal operation of the components. Therefore, the proof of the composition can be done *before* the components are proved or even implemented; all that is required is that the components have external specifications. Because of this characteristic of the proof methods, a software system can be specified, implemented and proved by starting at the top and working downward rather than by building upward from the Gypsy pre-defined functions and procedures. Thus, when working from the top down, the proofs provide a sound, scientific basis for predicting how the system will behave long before it can be run. It is in these high levels of system design that proofs often can be most effective.

3. The Gypsy environment

The Gypsy verification environment is an interactive program that supports a software engineer in specifying, implementing and proving Gypsy software systems. The specific goals of the environment are to increase the productivity of the software engineer and to reduce the probability of human error. To meet these goals, the Gypsy environment provides an integrated collection of conventional software development tools along with special tools for constructing formal, mathematical proofs. Figure 2 shows the logical structure of the environment.

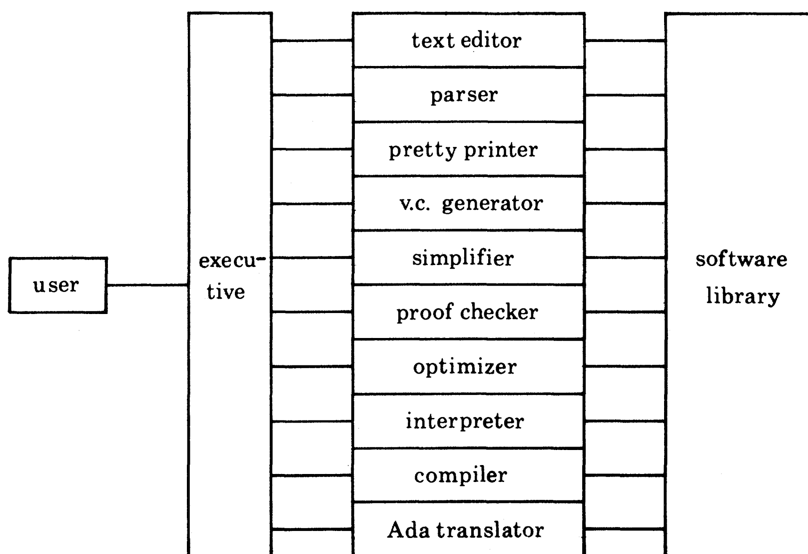

Figure 2. Gypsy environment components.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org

A single user interacts with the executive component of the environment to use a number of different software tools to build and evolve a software library. This library contains the various Gypsy components of the specification and implementation of a software system, as well as other supporting information such as verification conditions and proofs. The executive notes the changes that are made as the library evolves and marks components that need to be reconsidered to preserve the validity of the proofs (Moriconi 1977).

The Emacs text editor (Stallman 1980), parser and pretty printer are conventional tools for creating and modifying Gypsy text. The parser transforms Gypsy text into an internal form for storage in the library. The pretty printer transforms the internal form back into parsable Gypsy text. The interpreter, compiler (Smith 1980) and Ada translator (Akers 1983) also are fairly conventional tools for running Gypsy programs. Although the interpreter would be a very useful debugging tool, it is not well developed and it is not presently available.

The tools that are involved in constructing proofs are the verification condition generator, the algebraic simplifier, the interactive proof checker and the optimizer. The verification condition generator automatically constructs verification conditions from the Gypsy text of a program. The algebraic simplifier automatically applies an *ad hoc* set of rewrite rules that reduce the complexity of the verification conditions and other logical formulas produced within the Gypsy environment. These rewrite rules are based on equality (and other) relations that are applied by the definitions of the Gypsy pre-defined functions. The interactive proof checker has evolved from one described by Bledsoe & Bruell (1974). It provides a set of truth preserving transformations that can be performed on first-order predicate calculus formulas. These transformations are selected interactively.

The optimizer (McHugh 1983) is unique to the Gypsy environment. It produces logical formulas whose truth is sufficient to show that certain program optimizations are valid. The optimizer works in a manner similar to the verification condition generator. From the implementation of a program *and* its specifications, logical formulas called *optimization conditions* are constructed automatically. These conditions are proved, and then the compiler uses this knowledge to make various optimizations.

4. An example

To illustrate the capabilities of the Gypsy language and environment, consider the design of a simple software system that filters a stream of messages. Two computers, A and B, are to be coupled by a transmission line so that A can send messages to B. These messages are strings of ASCII characters arranged in a certain format. However, certain kinds of these messages, even when properly formatted, cause machine B to crash. To solve this problem a separate microcomputer is to be installed between A and B as shown in figure 3. The microcomputer is to monitor the flow of messages from A to B, remove the undesirable messages and log them on an audit trail.

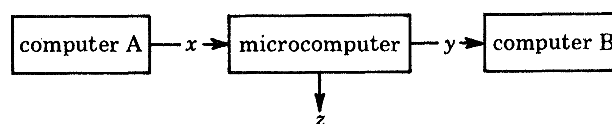

Figure 3. Microcomputer filter.

Downloaded from rsta.royalsocietypublishing.org

(a) *Top level specification*

The microcomputer filter will be developed from the top downwards. The process begins by defining an abstract specification of its desired behaviour. The Gypsy text for this top level specification is shown in figure 4. When using the Gypsy environment, the first step would be to create this text and store it in the software library.

```
scope message_stream_separator =
begin

  procedure separator(x:a_char_seq; var y, z:a_char_seq) =
  begin
  exit separated(msg_stream(x), y, z);
    pending;
  end;

  function msg_stream(x:a_char_seq):a_msg_seq = pending;

  function separated(s:a_msg_seq; y, z:a_char_seq):boolean =
  pending;

  type a_char_seq = sequence of character;
  type a_msg = a_char_seq;
  type a_msg_seq = sequence of a_msg;

end;
```

FIGURE 4. Microcomputer filter top level specification.

The Gypsy text defines a scope called *message_stream_separator* that contains six Gypsy units, procedure *separator*, functions *msg_stream* and *separated* and types *a_char_seq*, *a_msg* and *a_msg_seq*. (A Gypsy scope is just a name that identifies a particular collection of Gypsy units. The Gypsy units are procedures, functions, constants, lemmas and types. All Gypsy programs are implemented and specified in terms of these five kinds of units.)

Procedure *separator* is the program that will filter the messages going from computer A to B. The external environment specification of *separator* is (*x*:*a_char_seq*; *var y*, *z*:*a_char_seq*). It states that *separator* has access to exactly three external data objects, x, y and z, as illustrated in figure 3. The object x is a constant, and y and z are variables. Each of the objects has a value that is a sequence of ASCII characters.

The operational specification is *exit separated*(*msg_stream*(x), y, z). This defines a relation among x, y and z that must be satisfied whenever *separator* halts (exits). The messages that arrive from computer A are supposed to be in a given format. However, there is no way to force A to deliver them properly, and even if it does there is the possibility of noise on the transmission line. Therefore, *separator* must be designed to extract properly formatted messages from an arbitrary sequence of characters. *Msg_stream*(x) is the function that applies the formatting rules and determines the sequence of properly formatted messages that are contained in an arbitrary sequence of characters. *Separated*(s, y, z) defines what it means for a sequence of messages s to be separated into two character strings y and z.

This top level specification does not give precise definitions for *msg_stream* and *separated*; only environment specifications for them are given. (The environment specifications for functions are interpreted in the same way as for procedures, except that the additional type name immediately preceding the '=' identifies the type of value produced by the function.) The

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org

precise definitions of *msg_stream* and *separated*, as well as the implementation of *separator*, are left *pending* at this stage of development. At this stage, the interface between *separator* and its external environment has been defined, and it has been acknowledged that *separator* must be prepared to deal with an input sequence that may contain improperly formed messages. Formulation of precise definitions for the pending items will be deferred to a later stage.

(*b*) *Specification refinement*

The next stage is to refine the operational specifications of *separator*. Figure 5 shows the actual Gypsy text that would be entered into the software library. This text extends scope *message_stream_separator* by replacing the old version of *separated* by the new one and by defining some new functions, types and lemmas.

In this refinement, the *separated* specification is given a precise definition in terms of two new functions *passed* and *rejected*. The definition is given by the operational specification of *separated*.

```
$extending
scope message_stream_separator =
begin

  function separated(s:a_msg_seq; y, z:a_char_seq):boolean =
  begin
  exit [assume result iff y = passed(s) & z = rejected(s)];
  end;

  function passed(s:a_msg_seq):a_char_seq =
  begin
  exit [assume result =
        if s = null(a_msg_seq) then null(a_char_seq)
        else passed(nonlast(s)) @ image(last(s)).pass fi];
  end;

  function rejected(s:a_msg_seq):a_char_seq =
  begin
  exit [assume result =
        if s = null(a_msg_seq) then null(a_char_seq)
        else rejected(nonlast(s)) @ image(last(s)).reject fi];
  end;

  function image(m:a_msg):an_image = pending;

  type an_image = record(pass, reject:a_char_seq);

  lemma null_separation =
  separated(null(a_msg_seq), null(a_char_seq),
                             null(a_char_seq));

  lemma extend_separation(s:a_msg_seq; m:a_msg;
                          y, z:a_char_seq) =
  separated(s, y, z)
  -> separated(s @ [seq: m], y @ image(m).pass,
                             z @ image(m).reject);

  lemma null_stream =
  msg_stream(null(a_char_seq)) = null(a_msg_seq);

end;
```

FIGURE 5. Microcomputer filter specification refinement.

Downloaded from rsta.royalsocietypublishing.org

Result is the Gypsy convention for the name of the value returned by a function, and the specification states that *result* is to be true iff $y = passed(s)$ and $z = rejected(s)$. The keyword *assume* indicates that this specification is to be assumed without proof. This is the normal Gypsy style for defining a function that is to be used just for specification.

Functions *passed* and *rejected* are defined in terms of pre-defined Gypsy functions and the function *image*. *Last* is a pre-defined function that gives the last element of a non-empty sequence, and *nonlast* gives all the other elements. The operator '@' denotes a pre-defined function that appends two sequences.

Image is a function that takes a message and produces a record of two parts, *pass* and *reject*. At a subsequent development stage, the definition of *image* will be refined to include the criterion for identifying a message that causes computer B to crash. *Image* also will define the actual output that is sent to computer B *and* to the audit trail for *each* message. If the message is of the form that will cause B to crash, the *pass* part of the record will contain a null sequence of characters and the *reject* part will contain the offending message and any other appropriate information. This record form for the result of *image* was chosen so that messages that are forwarded to B also can be audited if desired. This can be done by sending characters to both the *pass* and *reject* parts of the record. This design choice retains a large amount of flexibility for the subsequent design of the audit trail. The function *passed* applies the *image* function to each successive message m and appends the *pass* part of $image(m)$ to y. Similarly, *rejected* applies *image* to each m and appends the *reject* part to z.

(c) *Specification proof*

The Gypsy text for the specification refinement also contains three lemmas. These are properties that can be proved to follow from the preceding definitions. These lemmas are the beginning of a simple problem domain theory of separating messages. The lemmas (theorems) of this theory serve several important purposes. First, to the extent that they are properties that the software designer intuitively believes *should* follow from the assumed definitions, proving that they *do* follow provides confidence in these assumptions. Secondly, these properties are the basis for the implementation in the next stage. They are used in the proof of the implementation to decompose the proof into manageable parts. Thirdly, to the extent that the lemmas in this theory are reusable, they can significantly reduce the cost of other proofs that are based on the same theory (Good 1982*a*).

The *null_separation* lemma is a rather trivial one that states that if a sequence of messages s is empty, then $separated(s, y, z)$ is satisfied if y and z also are empty. Lemma *extend_separation* describes how to extend the *separated* relation to cover one more message m. If $separated(s, y, z)$ is satisfied, then so is $separated(s@[seq:m], y@image(m).pass, z@image(m).reject)$.

A formal proof of both of these lemmas can be constructed with the assistance of the interactive proof checker in the Gypsy verification environment. The proof checker provides a fixed set of truth-preserving transformations that can be performed on a logical formula. Although the proof checker has some very limited capability to make transformations without user direction, the primary means of constructing a proof is for the user to select each transformation. Expansion of the definition of a function is one kind of transformation that can be made. The user directs the proof checker to expand the definition of a particular function, and then the expansion is done automatically. Other examples of transformations provided by the proof checker are instantiation of a quantified variable, substitution of equals for equals,

Downloaded from rsta.royalsocietypublishing.org

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

THE ROYAL SOCIETY A

PHILOSOPHICAL TRANSACTIONS OF

and use of a particular lemma. A formula is proved to be a theorem by finding a sequence of transformations that transform the formula into *true*. This sequence constitutes a formal, mathematical proof.

A complete transcript of the interactive proof of *extend_separation* is given in Appendix A. The key steps in the proof are to expand the definition of the *separated* relation and the *passed* and *rejected* functions with the expand command. The theorem command shows the state of the formula at various intermediate stages of transformation. The *null_separation* lemma is proved in a similar way.

Notice that both of these lemmas about message separation can be proved at this rather high level of abstraction without detailed knowledge of the specific format for incoming messages and without knowing the specific formatting details for the outputs y and z. These details are encapsulated in the functions *msg_stream* and *image* respectively. These definitions (which would need to be provided in subsequent refinement stages) might be quite simple or very complex. In either case, however, detailed definitions of these functions are *not* required at this stage. The use of abstraction in this way is what makes it possible to construct concise, intellectually manageable, formal proofs about large complex specifications. The next §4 (*d*) illustrates how similar techniques can be used in proofs about an implementation.

Finally, it is noted that the *null_stream* lemma cannot be proved at this stage of refinement. However, it is required in the subsequent implementation proof, and therefore, it serves as a constraint on the refinement of the definition of *msg_stream*.

(*d*) *Implementation refinement*

An implementation of procedure *separator* that satisfies the preceding specifications is shown in figure 6. The implementation contains two internal variable objects m and p of types *a_msg* and *integer* respectively. *Separator* causes its effect on its external variable objects, y and z, first by assigning each of them the value of the empty sequence of characters; then it enters a loop that separates the messages in x one by one, and for each message the appropriate output is appended to y and z.

The desired effect of the loop is described by the *assert* statement. It states that on each iteration of the loop, messages in the subsequence $x[1..p]$ have been separated. (The Gypsy notation for element i of sequence x is $x[i]$, and $x[1..p]$ is the notation for the subsequence $x[1], \dots, x[p]$.) This assertion is an *internal* specification about the operation of the procedure.

The loop operates by successively calling the procedures *get_msg* and *put_msg*. *Get_msg* assigns to m the next properly formatted message in x and increases p to be the number of the last character in x that has been examined. *Put_msg* appends to y and z the appropriate output for the new message m. The properties of *get_msg* and *put_msg* are stated precisely in the specifications that are given for them in figure 6. (For the variable p, p' refers to its value at the time *get_msg* is started running, and p refers to its value when the procedure halts. The operator <: appends a single element to the end of a sequence.)

(*e*) *Implementation proof*

The remaining task for this level of the design of the microcomputer filter is to prove that this abstract implementation of *separator* satisfies its specifications (both internal and external). This proof is possible without any further refinement of the specifications or the implementation. The current form is an instance of the one shown in figure 1. Specifications and an

Downloaded from rsta.royalsocietypublishing.org


```
$extending
scope message_stream_separator =
begin

procedure separator(x:a_char_seq; var y, z:a_char_seq) =
begin
exit separated(msg_stream(x), y, z);
  var m:a_msg;
  var p:integer := 0;
  y := null(a_char_seq);
  z := null(a_char_seq);
  loop assert separated(msg_stream(x[1..p]), y, z)
       & p le size(x);
    if p = size(x) then leave;
    else get_msg(x, m, p);
         put_msg(m, y, z);
    end;
  end;
end;

procedure get_msg(x:a_char_seq; var m:a_msg; var p:integer) =
begin
exit msg_stream(x[1..p]) = msg_stream(x[1..p']) <: m
     & p > p' & p le size(x);
  pending
end;

procedure put_msg(m:a_msg; var y, z:a_char_seq) =
begin
exit y = y' @ image(m).pass & z = z' @ image(m).reject;
  pending
 end;

end;
```

FIGURE 6. Microcomputer filter implementation refinement.

implementation for *separator* have been constructed, but there is no implementation of either *get_msg* or *put_msg*. This level of proof simply assumes that these procedures eventually will be implemented and proved to satisfy their specifications. However, at this level, only their external specifications are required.

It is easy to see that the *exit* specification of *separator* logically follows from the assert statement in the loop whenever the procedure leaves the loop. This follows simply from the facts that, when the loop halts, $p = size(x)$ and that for every Gypsy sequence $x[1\,..\,size(x)] = x$. It is also easy to see that the assert statement is true the first time the loop is entered. This is because the local variable p is zero, and y and z are both equal to the empty sequence. The assertion then follows from the *null_stream* and *null_separation* lemmas because in Gypsy, $x[1\,..\,0]$ is the empty sequence and the size of a sequence is always non-negative. Finally, the *extend_separation* lemma can be used to prove that if the loop assertion is true on one iteration of the loop, then it also is true on the next. These steps form an inductive proof that the loop assertion is true on every iteration of the loop (even if it never halts). The loop, however, does halt because, according to the specifications of *get_msg*, p is an integer that increases on each iteration and yet never increases beyond the number of characters in the constant x. Therefore, the loop must halt; and when it does, the *exit* specification follows from the loop assertion.

Downloaded from rsta.royalsocietypublishing.org

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

The Gypsy verification environment automates all of this argument (except the argument about the loop halting). From the Gypsy text shown in figure 6, the verification conditions generator automatically constructs the formulas shown in figure 7.

```
Verification condition separator#2
separated (msg_stream (null (#seqtype#)),
           null (a_char_seq), null (a_char_seq))

Verification condition separator#3
  H1: msg_stream (x[1..p]) @ [seq: m#1] = msg_stream (x[1..p#1])
  H2: y @ image (m#1).pass = y#1
  H3: z @ image (m#1).reject = z#1
  H4: separated (msg_stream (x[1..p]), y, z)
  H5: p le size (x)
  H6: p + 1 le p#1
  H7: p#1 le size (x)
  H8: size (x) ne p
 -->
  C1: separated (msg_stream (x[1..p#1]), y#1, z#1)
```

FIGURE 7. Separator verification conditions.

Verification condition *separator#2* is the formula that states that the loop assertion is true the first time the loop is entered. *Separator#3* is the one that states that if the assertion is true on one iteration of the loop, it also is true on the next. Lines labelled *Hi* are the hypotheses of an implication, and lines labelled *Ci* are conclusions. Both the hypotheses and the conclusions are connected implicitly by logical conjunction. The notation *m#1* denotes a value of *m* upon completing the next cycle of the loop, and similarly for *p*, *y* and *z*. The notation [*seq*: *m#1*] means the sequence consisting of the single element *m#1*. The verification condition generator also has constructed a *separator#4* for the case when the loop terminates. The generator, however, does not present this one because the formula has been proved automatically by the algebraic simplifier. The best way to see the effect of the simplifier is to see what the verification conditions look like without it. The unsimplified formulas are shown in figure 8. (There also is a *separator#1*, which is so trivial that the generator does not even bother to use the algebraic simplifier.)

A complete transcript of the interactive proof of *separator#3* is given in Appendix B. The key steps are to do equality substitutions based on hypotheses H1, H2 and H3 with the eqsub command and then use the *extend_separation* lemma. *Separator#2* is proved by use of the lemmas *null_stream* and *null_separation*.

Once *separator* has been proved, the process of refinement can be resumed. In general, the refinement of both specifications and implementations is repeated until all specifications and procedures are implemented in terms of Gypsy primitives.

It is important to observe that the proof of *separator* has identified formal specifications for *get_msg* and *put_msg* that are adequate for the subsequent refinements of these procedures. It has been proved that *separator* will run according to its specification if *get_msg* and *put_msg* run according to theirs. Therefore, these specifications are completely adequate constraints for the subsequent refinements. Some of the specifications may not be necessary, But they are sufficient to ensure that *separator* will satisfy its specification.

Downloaded from rsta.royalsocietypublishing.org


```
Verification condition separator#2
  H1: true
 -->
  C1: separated (msg_stream (x[1..0]), null (a_char_seq),
                 null (a_char_seq))
  C2: 0 le size (x)

Verification condition separator#3
  H1: separated (msg_stream (x[1..p]), y, z)
      & p le size (x)
  H2: not p = size (x)
  H3: msg_stream (x[1..p#1]) = msg_stream (x[1..p]) <: m#1
      & p#1 > p
      & p#1 le size (x)
  H4: y#1 = y @ image (m#1).pass
      & z#1 = z @ image (m#1).reject
 -->
  C1: separated (msg_stream (x[1..p#1]), y#1, z#1)
  C2: p#1 le size (x)

Verification condition separator#4
  H1: separated (msg_stream (x[1..p]), y, z)
      & p le size (x)
  H2: p = size (x)
 -->
  C1: true
  C2: separated (msg_stream (x), y, z)
```

FIGURE 8. Unsimplified verification conditions.

5. TRIAL APPLICATIONS

The Gypsy environment has been developed to explore the practicality of constructing formal proofs about software systems that are intended to be used in actual operation. Throughout its development, the environment has been tested on a number of trial applications. The two major ones are summarized in §5(*a*), (*b*).

(*a*) *Message flow modulator*

The most recent application of Gypsy is the message flow modulator (Good *et al.* 1982*b*). The microcomputer filter that has been specified, designed and proved in §4 is a good approximation of the modulator. The microcomputer filter example was chosen deliberately to show how it is possible to construct concise, formal proofs about much larger software systems. The modulator consists of 556 lines of implementation, and the proofs in the preceding sections apply, with only very minor alteration, to the design of the modulator. The lower level details that are unique to the modulator are encapsulated in the *msg_stream* and *image* functions.

The message flow modulator is a filter that is applied continuously to a stream of messages flowing from one computer system to another. As in the microcomputer filter, messages that pass the filter are passed on to their destination with a very minor modification. Messages that do not are rejected and logged on an audit trail. A properly formatted message consists of a sequence of at most 7200 ASCII characters that are opened and closed by a specific sequence.

The filter consists of a list of patterns. Each pattern defines a sequence of letters and digits that may be interspersed with various arrangements of delimiters (a delimiter is any character

Downloaded from rsta.royalsocietypublishing.org

other than a letter or digit). If a message contains any phrase that matches any pattern, it is rejected to the audit trail along with a description of the offending pattern. Messages that do not contain any occurrence of any pattern are forwarded on to their destination.

In essence, the formal specifications of the modulator have the form $y = f(x, r)$ & $z = g(x, r)$, where r is the list of rejection patterns. The specification describes the exact sequences of characters that must flow out of the modulator for every possible input sequence. This includes handling both properly and improperly formatted messages in the input stream, detecting phrases that match the rejection patterns, and formatting both output sequences. The Gypsy formulation of these specifications is described in further detail in Good *et al.* (1982*b*).

The modulator was developed within the Gypsy environment as a converging sequence of prototypes. First, Gypsy specifications and proofs were constructed for the top levels of the modulator design. This design covered the basic separation of messages into the two output streams. Then, a sequence of running prototypes was implemented. The purpose of these prototypes was to help decide what some of the detailed behaviour of the modulator *should* be. These prototypes were used to investigate various approaches to handling improperly formed messages and to formatting the audit trail. Specifications for these aspects of the modulator were decided upon only after considerable experimentation with the prototypes. Next, another sequence of performance prototypes was built to evaluate the performance of various pattern matching implementations. Once adequate performance was attained, the Gypsy specifications and proofs were completed for the entire modulator.

As the final step, the proved modulator was tested in a live, operational environment on test scenarios developed by an independent, external group. Without any modification, the proved modulator passed all of these tests on the first attempt.

(*b*) *Network interface*

The first major application of Gypsy, and the most complex one to date, was a special interface for the ARPANET. Each ARPANET host has message traffic that needs to be transported over the network according to the standard Transmission Control Protocol (Version 4.0). The ARPANET, however, is assumed to be an untrustworthy courier. The special interfaces are to ensure proper message delivery across this potentially unreliable network.

Normally, each host is connected directly to the network by a bi-directional cable. Each cable is cut and an interface unit is installed at the cut (figure 9). This turns the 'dumb' cable into

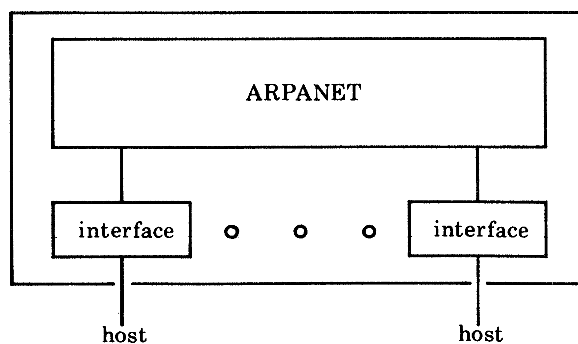

FIGURE 9. ARPANET interface.

Downloaded from rsta.royalsocietypublishing.org

a 'smart' one. When the smart cable receives a message from the host, the message is checked to see that it is return-addressed to the sending host. If it is not, the message is dropped. If it is properly return-addressed, then, in effect, the smart cable seals the message in a plain brown envelope that cannot be opened by the network, addresses and return-addresses the envelope and sends it to the ARPANET for delivery. In the other direction, when the cable receives an envelope from the network, it is opened and its message is examined. If the message shows no sign of damage and it is addressed to the receiving host, the message is forwarded on to the host; otherwise it is dropped. So if the network behaves in some unpredictable way and delivers an envelope to the wrong smart cable, the smart cable detects this and refuses to deliver the message to its host.

The specification for the interface unit is relatively straightforward. It states that all messages that are sent into the network must be properly return addressed and packaged and that all messages that are sent out to the host must be properly unpackaged and addressed. The implementation of the interface, however, is rather involved because of a variety of fixed, external constraints. One of the major constraints was that the interface was required to use standard ARPANET link and transport protocols; another was the hardware that was to run the interface. The interface hardware consisted of two PDP-11 minicomputers packaged into a single box. One PDP-11 was connected to the host, and the other was connected to the network. The two PDP-11s could communicate only in an extremely restricted way, and no other communication between the host and the network was allowed.

The proved network interface also was developed as a converging sequence of prototypes. First, the formal specification and proof methods were applied at the highest levels of system design. This involved a specification and proof about the concurrent host-to-host communication across the entire set of interfaces (including the ARPANET). Then, the formal specification and proof methods were applied to obtain the distribution of the interface software onto the two PDP-11 processors of the actual hardware. At this point, a sequence of running prototypes was implemented to evaluate the performance that could be attained with this design. The resources required by the initial design were much greater than those that were available, so a new design was developed and proved, and new performance prototypes were built. When adequate performance was attained, the formal specification and proof methods were applied through all the remaining levels of the interface design and implementation. The general approach that was used on the network interface is illustrated in Good (1982*c*).

The final result was a formally specified and proved interface, implemented in Gypsy, that operated successfully across the ARPANET with a companion interface that was implemented independently, in conventional assembly code, by Bolt, Beranek and Newman, Inc. As in the flow modulator, without any modification, the proved interface worked properly in every trial run. (A small number of inadequacies in the statement of the constraints for the message formats and protocols were detected and fixed during the prototype stage.)

(*c*) *Economics*

These trial applications indicate the kinds of specifications and proofs that are within the capability of the Gypsy environment. However, if formal specifications and proof are to be used as the basis for a new approach to software engineering, there also is the matter of economics. Table 1 shows various measures of the scale of the trial applications and estimates of the amounts of resources used.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES
PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A

Table 1. Resources used

	flow modulator	network interface
lines of Gypsy specifications	1283	3135
lines of executable Gypsy	556	4211
words of compiled PDP-11 code	3849	42271
verification conditions	304	2600
lemmas used	44	90
automatic proofs	146	2287
interactive proofs mechanically checked	198	313
lemmas assumed without proof	4	2
work-months	13	52
DEC 2060 c.p.u.-hours	220	444
page-months of file storage	45000	84465
proved, total Gypsy lines per work-day	6.43	6.42
proved, total Gypsy lines per c.p.u.-hour	8.36	16.54
proved executable Gypsy lines per work-day	1.94	3.68
proved, executable Gypsy lines per c.p.u.-hour	2.53	9.48

The 'lines of Gypsy specifications' and 'lines of executable Gypsy' must be interpreted with caution. These count actual lines of Gypsy text (excluding comments). A line count, however, obviously depends on the style in which the Gypsy text is written, and therefore these counts are quite subjective. Also, a line count is not necessarily a good measure of complexity. In spite of these obvious weaknesses, line counts are one of the most frequently quoted measures of program size.

'Lines of Gypsy specifications' refers to those lines that are used to express formal specifications. One of the important differences between the flow modulator and the network interface was the strength of their specifications. The specifications for the flow modulator were very strong; they completely defined the two output sequences as functions of the input sequence. The specifications for the network interface, however, were much weaker; they were stated as relations rather than as functions, and there are some important aspects of the behaviour of the interface that are not covered by these relations. The difference between these two specification forms is like the difference between $y = f(x)$ and $y < x$. The first defines exactly what y must be for every x and the second states only a relatively weak relation between x and y. This is an important difference to consider in interpreting the numbers in table 1.

'Lines of executable Gypsy' refers to lines of Gypsy that actually cause run-time code to be compiled. These line counts do not, for example, include type declarations. 'Words of compiled PDP-11 code' refer to the number of (16 bit) words of PDP-11 machine code that were produced by the Gypsy compiler. In both applications, the target machine was a PDP-11/03 with no operating system other than the Gypsy run-time support package. This package is not included in the word count. The two applications were compiled with different Gypsy compilers. The flow modulator was compiled through the Gypsy-to-Bliss translator (Smith 1980). The network interface was compiled with the original Gypsy compiler (Hunter 1981).

'Verification conditions' refers to the number of verification conditions constructed by the Gypsy environment. 'Lemmas used' refers to the number of these stated in the Gypsy text. 'Automatic proofs' refers to the number of verification conditions proved fully automatically

Downloaded from rsta.royalsocietypublishing.org

by the algebraic simplifier. 'Interactive proofs mechanically checked' refers to the number of verification conditions and lemmas that required the use of the proof checker. In both applications, a small number of lemmas were assumed without proof. The four lemmas that were not proved at the time the flow modulator was completed have since been proved. The two lemmas that were not proved for the network interface were key assumptions about the problem domain.

'Work-months' refers to the total number of (22 day) working months required to complete the application. These months include the full development of the application, from initial conception through the final testing of the proved software. This includes all iterations of all levels of specifications, prototypes and proofs. Similarly, 'c.p.u.-hours' and 'page-months' also cover the full development cycle. 'Proved, total Gypsy lines' is computed from 'lines of Gypsy specifications' plus 'lines of executable Gypsy'. This gives a measure of the total number of Gypsy lines produced per working day. 'Proved, executable Gypsy lines' considers just 'lines of executable Gypsy'.

6. Conclusion

The Gypsy verification environment is an experimental system that has been developed to explore the practicality of a new approach to software engineering that is based on rigorous, mathematical foundations. These foundations, together with the tools provided in the Gypsy environment, make it possible for a software engineer to construct formal, mathematical proofs about a software system. By appropriate use of abstraction, the formal proofs can be kept concise and intellectually manageable even though they cover large, complex systems. These proofs provide an objective, scientific basis for predicting, accurately and precisely, how a software system will behave when it runs. These proofs can be constructed at all stages of the software life cycle, from the earliest design stages through to system maintenance. Therefore, they also provide the software engineer a basis for evaluating the effects of early design decisions at the time they are made rather than having first to build a system that runs. The proofs also provide a basis for predicting the effects of maintenace modifications.

The results of the first trial applications of the Gypsy environment have been very encouraging. The flow modulator and the network interface are non-trivial software systems. They are intended to be used in actual operation, and their predictability is a genuine, major concern. Although these applications do not approach the scale and complexity of what normally are regarded as 'large' systems, they do support the claim that a formal, mathematical approach to software engineering is technically viable. The next major research goal seems to be making this approach economically viable. Although the cost of applying this new technology in the two applications was much less than what might have been expected (and one always must weigh the cost of applying this mathematical approach against the cost of an unpredictable software system), there seem to be many ways in which the amount of resources used to apply the technology can be reduced. If this can be done, this new technology can become the basis for a new practice of software engineering that can provide dramatic improvements in the predictability and quality of software systems.

On this euphoric note, it is all too easy to be lulled into a false sense of security because it is tempting to believe that a formally specified and proved program should be absolutely correct. It should always behave perfectly and never malfunction. However, there are several

Downloaded from rsta.royalsocietypublishing.org

reasons why a program that has been proved within the Gypsy environment may not behave exactly as expected. First, the formal specifications may not describe exactly the expected behaviour of the program. Secondly, the formal specifications may not describe all of the aspects of program behaviour. Thirdly, invalid lemmas may have been assumed without proof. Finally, either the verification environment, the compiler, the Gypsy run-time support or the hardware might malfunction.

The last of these potential sources of error, in principle, can be minimized by specifying and proving the verification environment, the compiler, the run-time support and to some degree, the hardware. These would be large complex proofs that are well beyond present capabilities; but, given sufficient cost reductions, these proofs eventually may well be possible. The first three, however, are subjective and involve some element of human judgment. Therefore, these potential sources of error cannot be eliminated. These sources of error are cited not to belittle the potential of a scientific basis for software engineering but to make clear that the formal, mathematical approach offers no absolutes. As with any other science, it must be applied in the context of human judgment.

The development and initial experimental applications of Gypsy have been sponsored in part by the U.S. Department of Defense Computer Security Center (contracts MDA904-80-C-0481, MDA904-82-C-0445), by the U.S. Naval Electronic Systems Command (contract N00039-81-C-0074), by Digital Equipment Corporation, by Digicomp Research Corporation and by the National Science Foundation (grant MCS-8122039).

Appendix A. Formal proof of lemma *extend_separation*

The following is the complete transcript of the interactive proof of the lemma *extend_separation*. The input supplied by the human user is underlined.

```
Entering Prover with lemma extend_separation

   H1: separated (s, y, z)
->
   C1: separated (s @ [seq: m],
                  y @ image (m).pass,
                  z @ image (m).reject)

Prvr -> expand
       Unit name -> separated
Which ones?
   1. in H1:  separated (s, y, z)
   2. in C1:  separated (s @ [seq: m],
                         y @ image (m).pass,
                         z @ image (m).reject)

          <number-list>, ALL, NONE, PRINT, ^E: all
```

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org


```
Prvr -> theorem
   H1: passed (s) = y
   H2: rejected (s) = z
->
   C1: y @ image (m).pass = passed (s @ [seq: m])
   C2: z @ image (m).reject = rejected (s @ [seq: m])

Prvr -> expand
       Unit name -> passed
Which ones?
   1. in H1:  passed (s)
   2. in C1:  passed (s @ [seq: m])

        <number-list>, ALL, NONE, PRINT, ^E: 2

Prvr -> expand
       Unit name -> rejected
Which ones?
   1. in H2:  rejected (s)
   2. in C2:  rejected (s @ [seq: m])

        <number-list>, ALL, NONE, PRINT, ^E: 2

Prvr -> theorem
   H1: passed (s) = y
   H2: rejected (s) = z
->
   C1: y @ image (m).pass = passed (s) @ image (m).pass
   C2: z @ image (m).reject = rejected (s) @ image (m).reject

Prvr -> qed

9.  ANDSPLIT
   11.  SIMPLIFYC
      14.  UNIFY
   12.  SIMPLIFYC
      18.  UNIFY
Theorem proved!
```

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Downloaded from rsta.royalsocietypublishing.org

Appendix B. Formal proof of verification condition *separator#3*

The following is the complete transcript of the interactive proof of the verification condition *separator#3*. The input supplied by the human user is underlined.

```
Entering Prover with verification condition separator#3

   H1: msg_stream (x[1..p]) @ [seq: m#1] = msg_stream (x[1..p#1])
   H2: y @ image (m#1).pass = y#1
   H3: z @ image (m#1).reject = z#1
   H4: separated (msg_stream (x[1..p]), y, z)
   H5: p le size (x)
   H6: p + 1 le p#1
   H7: p#1 le size (x)
   H8: size (x) ne p
->
   C1: separated (msg_stream (x[1..p#1]), y#1, z#1)

Prvr -> retain
       hypothesis labels, ALL, NONE -> h1 h2 h3 h4

Prvr -> theorem
   H1: msg_stream (x[1..p]) @ [seq: m#1] = msg_stream (x[1..p#1])
   H2: y @ image (m#1).pass = y#1
   H3: z @ image (m#1).reject = z#1
   H4: separated (msg_stream (x[1..p]), y, z)
->
   C1: separated (msg_stream (x[1..p#1]), y#1, z#1)

Prvr -> eqsub
        Hypothesis label -> h1
  msg_stream (x[1..p#1]) := msg_stream (x[1..p]) @ [seq: m#1]

Prvr -> theorem
   H1: y @ image (m#1).pass = y#1
   H2: z @ image (m#1).reject = z#1
   H3: separated (msg_stream (x[1..p]), y, z)
->
   C1: separated (msg_stream (x[1..p]) @ [seq: m#1], y#1, z#1)

Prvr -> eqsub
        Hypothesis label -> h1
  y#1 := y @ image (m#1).pass

Prvr -> theorem
   H1: z @ image (m#1).reject = z#1
   H2: separated (msg_stream (x[1..p]), y, z)
->
   C1: separated (msg_stream (x[1..p]) @ [seq: m#1],
                  y @ image (m#1).pass, z#1)
```

Downloaded from rsta.royalsocietypublishing.org


```
Prvr -> eqsub
         Hypothesis label -> h1
  z#1  := z @ image (m#1).reject

Prvr -> theorem
   H1: separated (msg_stream (x[1..p]), y, z)
->
   C1: separated (msg_stream (x[1..p]) @ [seq: m#1],
                  y @ image (m#1).pass,
                  z @ image (m#1).reject)

Prvr -> use
        Unit name -> extend separation

Prvr -> theorem
   H1:    separated (s$#2, y$#2, z$#2)
       -> separated (s$#2 @ [seq: m$#2],
                     y$#2 @ image (m$#2).pass,
                     z$#2 @ image (m$#2).reject)
   H2: separated (msg_stream (x[1..p]), y, z)
->
   C1: separated (msg_stream (x[1..p]) @ [seq: m#1],
                  y @ image (m#1).pass,
                  z @ image (m#1).reject)

Prvr -> proceed

11.  BACKCHAIN
   12.  UNIFY
      13.  ANDSPLIT
   13.  UNIFY
Theorem proved!.
```

References

Akers, R. L. 1983 A Gypsy-to-Ada program compiler. Master's thesis, University of Texas at Austin. Also *Tech. Rep.* no. 39, Institute for Computing Science, The University of Texas at Austin.

Bledsoe, W. W. & Bruell, P. 1974 A man–machine theorem-proving system. In *Advance Papers of Third International Joint Conference on Artificial Intelligence* (ed. W. W. Bledsoe), **5–1** (Spring), pp. 51–72.

Dijkstra, E. W. 1968 A constructive approach to the problem of program correctness. *BIT* **8**, 174–186.

Dijkstra, E. W. 1976 *A discipline of programming*. Englewood Cliffs, N.J.: Prentice-Hall.

Floyd, R. W. 1967 Assigning meanings to programs. In *Proceedings of a Symposium in Applied Mathematics* (ed. J. T. Schwartz), vol. 19, pp. 19–32. Providence, Rhode Island: American Mathematical Society.

Good, D. I. 1970 Toward a man–machine system for proving program correctness. Ph.D. thesis, University of Wisconsin.

Good, D. I., Cohen, R. M., Hoch, C. G., Hunter, L. W. & Hare, D. F. 1978 Report on the language Gypsy, Version 2.0. *Tech. Rep.* ICSCA-CMP-10, Certifiable Minicomputer, Project, ICSCA, The University of Texas at Austin.

Good, D. I. 1982*a* Reusable problem domain theories. In *Formal Specification – Proceedings of the Joint IBM/University of Newcastle-upon-Tyne Seminar* (ed. M. J. Elphick), pp. 92–115. Also *Tech. Rep.* no. 31, Institute for Computing Science, The University of Texas at Austin.

Good, D. I., Siebert, Ann E. & Smith, L. M. 1982*b* Message Flow Modulator – Final Report. *Tech. Rep.* no. 34, Institute for Computing Science, The University of Texas at Austin.

Downloaded from rsta.royalsocietypublishing.org

Good, D. I. 1982*c* The proof of a distributed system in Gypsy. In *Formal Specification – Proceedings of the Joint IBM/University of Newcastle-upon-Tyne Seminar* (ed. M. J. Elphick), pp. 443–489. Also *Tech. Rep.* no. 30, Institute for Computing Science, The University of Texas at Austin.
Gries, D. 1981 *The science of computer programming*. New York: Springer-Verlag.
Hoare, C. A. R. 1969 An axiomatic basis for computer programming. *Commun. Ass. comput. Mach.* **12–10**, 576–580.
Hoare, C. A. R. 1982 Programming is an engineering profession. *Tech. Rep.* no. PRG-27, Programming Research Group, Oxford University Computing Laboratory.
Hunter, L. W. 1981 *The first generation Gypsy compiler. Tech. Rep.* no. 23, Institute for Computing Science, The University of Texas at Austin.
Jensen, K. & Wirth, N. 1974 *Pascal user manual and report*. New York: Springer-Verlag.
Jones, C. B. 1980 *Software development: a rigorous approach*. Englewood Cliffs, N.J.: Prentice-Hall.
King, J. C. 1969 A program verifier. Ph.D. thesis. Carnegie–Mellon University.
McHugh, J. 1983 Toward the generation of efficient code from verified programs. Ph.D. thesis, University of Texas at Austin.
Moriconi, M. S. 1977 A system for incrementally designing and verifying programs. Also Ph.D. thesis, *Tech. Rep.* ICSCA-CMP-9. The University of Texas at Austin.
Naur, P. 1966 Proof of algorithms by general snapshots. *BIT* **6**, 310–316.
Neumann, P. G. 1983*a* Letters from the Editor. *Software engineering Notes* **8** (3), 2–6.
Neumann, P. G. 1983*b* Letters from the Editor. *Software engineering Notes* **8** (5), 1–9.
Smith, L. M. 1980 Compiling from the Gypsy verification environment. Master's thesis, The University of Texas at Austin. Also *Tech. Rep.* no. 20. Institute for Computing Science, The University of Texas at Austin.
Stallman, R. M. 1980 *EMACS Manual for Twenex Users*, M.I.T. Artificial Intelligence Laboratory.

Discussion

B. A. Wichmann (*National Physical Laboratory, Teddington, Middlesex, U.K.*). Could Dr Good comment upon the dependence of his work on the correctness of the compilers and hardware?

D. I. Good. Certainly, the claim that a proved program actually will run correctly is based on the assumption that it is compiled correctly and that the hardware runs correctly and on a number of other things. These are discussed briefly in the conclusion of the paper. This question, however, also raises two other important issues.

The first is program 'correctness'. This term is widely used, but I believe it is highly ill-chosen because it conveys a misleading connotation of absolute perfection. For example, when someone says they have proved that a program is 'totally correct', who would believe that it would ever do anything but always run absolutely perfectly? Yet, it is quite possible for such a program to malfunction. This is because when a program is proved, it is proved against a particular specification. The proof provides assurance that the program will run according to its specification; but, the proved program very well might do other things that are not covered by the specification, and it might do them wrong! Whenever someone claims to have proved a program, the first question should be 'what did you prove?'

The second question should be 'what did you assume?'. Every proof, whether it is about a program or anything else, is based upon certain assumptions. These may be very simple, or they may be arbitrarily complex. All that is produced in any proof is a chain of deductive steps that imply that a conclusion follows from a set of assumptions. If the assumptions are not true, then the conclusion need not be either. Thus, in the end, all that *any* proof does is to make explicit the assumptions upon which the conclusion is based. A proof provides us confidence in its conclusion to the extent that the conclusion is deduced from believable assumptions. For software systems, the simple accomplishment of identifying a precise set of assumptions that imply that a software system runs according to a particular specification is a dramatic improvement over conventional methods. There always remains, however, the question of the validity of the assumptions. For a proof to provide confidence that a system runs according to specification, its assumptions must be simple, concise, and believable.