

Termination Conventions and Comparative Semantics

Eric C.R. Hehner and Andrew J. Malton

University of Toronto, CSRI, Sandford Fleming Building, 10 King's College Road, Toronto,
Canada, M5S 1A4

Summary. The notion of termination is examined, first for its physical observability, then for its part in six semantic formalisms, with emphasis on predicative semantics.

Observing Termination

Imagine there is a computer in front of you right now. Perhaps it has a screen and keyboard, as is common. You can press the keys, and watch the patterns on the screen. Unfortunately, this computer was built and programmed by Tatooinians, whose language you do not know. The symbols on the keys are meaningless to you; the patterns on the screen are undecipherable. Fortunately, no-one is asking you what is being computed, or what the computation means. You are being asked a simpler question: When has the computation terminated? If you will simply report when the machine is finished, someone else will then interpret the results.

How do you know when termination occurs? If nothing has happened for five minutes, can you report termination? No; perhaps the machine is computing *Ackermann* (6,6), and still has a century to go. It is obvious that nontermination, or infinite looping, is not an observable event; without knowing the program, you can only guess that a computation is stuck in a loop. The purpose of this thought-experiment is to convince you that termination is also not an observable event.

How do you recognize termination on your own computer, programmed by someone who speaks your language? Probably you look for the word "done", or something similar, printed as the dying act. Or perhaps you look for a prompt from the operating system, following termination. It is possible, though not probable, that your computer has broken in mid-computation, and is now behaving erratically; it just printed the letters "done" (or the prompt) by chance as it malfunctions. Or perhaps the program was written by a joker, who thinks it funny to print "done" (or the prompt) in the middle of the compu-

tation, followed in a minute by "wait, there's more". You cannot be certain when a computation has finished even on your own machine, although the usual indicators are very reliable.

Predicative Semantics

A predicative formalism is a formalism for specifying and describing computations, including the semantics of programs, as predicates. One is invited to decide what observable quantities are of interest, and to express specifications as predicates having these quantities as free variables. A computation satisfies a specification if it instantiates the free variables in a way that satisfies the predicate.

A reasonable choice of "observable quantities of interest" is the communication sequence, or sequences, to and from a computer. In many computer systems and programming languages, there is something that indicates termination of a communication sequence. (In UNIX it is control-D; in Pascal it is the *eof* function.) But at a basic semantic level, this end-of-file indication is just another communication. It is only by agreement that we give it the special meaning "don't wait for more". We may decide to enshrine this meaning in a programming language's semantics; there are some advantages in doing so. Similarly, we can design into a language a communication with the meaning "please enter input now", and many other equally useful messages. Or instead, we may decide that the programmers, not the language designers, give meaning to the communications. We leave it to them to decide upon whether and how to prompt, to decide whether and how to say "this message is the last". The benefit is a simpler semantics.

A recent paper [3] presented and used a predicative formalism in which the observable quantities of interest are taken to be the initial values \hat{x}, \hat{y}, \dots and final values $\acute{x}, \acute{y}, \dots$ of some variables. Pointedly, termination and nontermination are not taken to be observable. It is left to the programmers to decide upon their termination convention. We assume knowledge of [3], or of its forerunner [2], and we now discuss the possible termination conventions.

Stability Conventions

A termination convention, or more generally a *stability convention*, is a predicate $\sigma(v)$ of the program variables (or state) v that is not identically true:

$$\neg \forall v. \sigma.$$

A state satisfying σ is called "stable", and a state not satisfying σ is called "unstable". (We disallow the trivial predicate "true" as a stability convention to avoid having to say "non-trivial stability convention" in the theorems and comparative semantics that follow.) Here are two examples.

Convention u (unstable state). Let us call one of the states u . Then, under this convention, we define $\sigma = (v \neq u)$. In words, u is taken to be an unstable state,

and all other states are considered stable. (The symbol \perp is sometimes used for a fictitious state reached after nontermination. In contrast, we intend the unstable state u to be an actual machine state; it is the state during a computation, after its start and before its termination. We are not considering u to be the bottom of a lattice.)

Convention b (busy bit). Let us call one boolean variable b . Then, under this convention, we define $\sigma = \neg b$. We might implement b as a light on the console. Before the computation has started, it is off (false); during the computation, it is on (true); when it goes off (false), the state is again stable and the computation has terminated.

From a coding viewpoint, convention u is the most economical; it wastes only one state to indicate instability, leaving all others to indicate results. Convention b makes half the states unstable; still, it wastes only one bit. From a recognition viewpoint, convention b is the most practical; we need to look at only one bit, not all the bits, to determine stability.

Respectfulness

In [3] we defined the “determined” operator ∇ for any specification $S(\dot{v}, \dot{v})$ as

$$\nabla S = \neg \forall \dot{v} \cdot S.$$

∇S is true of exactly those inputs \dot{v} for which we are interested in the output \dot{v} . We also considered a special case of the invariant construct

$$(\text{inv } P \cdot S) = (\dot{P} \wedge \nabla S \Rightarrow S \wedge \dot{P})$$

where P is any predicate of the program state. In these terms, we now define “respectfulness”.

A specification S is said to *respect* convention σ iff

$$\forall \dot{v} \cdot \forall \dot{v} \cdot S = (\text{inv } \sigma \cdot S).$$

The following six theorems afford us some insight into respectfulness.

Theorem 0. S respects σ iff

$$\forall \dot{v} \cdot \nabla S = \dot{\sigma} \wedge (\forall \dot{v} \cdot S \Rightarrow \dot{\sigma}).$$

For a determined specification, the intent is that the initial state should be stable, and any final state satisfying S should be stable. During the computation, the machine is in an unstable state. Suppose for a moment that we can observe the state during a computation. Under convention u , since there is only one unstable state, we cannot see *how* the computation is progressing – only *that* it is. Under convention b , we may be able to watch the sequence of unstable states, and see how the machine works. But our specifications do not speak of intermediate, unstable states; in that sense, they are not “of interest”.

Theorem 1. *Specification $S(\hat{v}, \hat{v})$ respects convention u iff it is both strict*

$$S(u, u)$$

and explosive

$$\forall \hat{v} \cdot (S(\hat{v}, u) \Rightarrow \forall \hat{v} \cdot S(\hat{v}, \hat{v})).$$

This means that from initial state u , final state u is possible; and if, from some initial state, final state u is possible, then, from that initial state, every final state is possible.

Theorem 2. *Specification S respects convention b iff it is both strict*

$$\forall \hat{v} \cdot \exists \hat{v} \cdot \hat{b} \Rightarrow S \wedge \hat{b}$$

and explosive

$$\forall \hat{v} \cdot (\exists \hat{v} \cdot S \wedge \hat{b}) \Rightarrow \forall \hat{v} \cdot S.$$

Again, "strictness" means that instability may lead to instability, and "explosiveness" means that if S allows the final state to be unstable, then it allows every final state.

Theorem 3. *Let S be an arbitrary specification. Then $\text{inv } \sigma \cdot S$ respects convention σ .*

All our programming connectives that were defined in [3] have the property that if their parts respect some convention, then the whole does.

Theorem 4. *If P and Q respect σ , then so do $(P \circ Q)$, $(P; Q)$, $(P \text{ or } Q)$ and $(\text{if } b \text{ then } P \text{ else } Q)$. If, assuming P respects σ , $S(P)$ does also, then $\text{loop } P: S(P)$ respects σ .*

It is especially nice that, under any stability convention, the difference between relational product and composition disappears.

Theorem 5. *If Q respects some convention σ , then*

$$P; Q = P \circ Q.$$

The Need for a Stability Convention

We must be able to decide when a computation is finished. Since termination is not an observable event, it is necessary to have and to respect a stability convention. But it is really of no interest which convention we choose. We would like to be able to write a simple specification that refers only to the desired results, not to the arbitrary convention. As a trivial example, we would like to write

$$\dot{x} = \dot{x} + 1$$

to say that x is increased by 1. Unfortunately, this specification does not respect any convention. To obtain a respectful specification, we apply Theorem 3. Under convention b we obtain

$$\neg \hat{b} \Rightarrow \dot{x} = \dot{x} + 1 \wedge \neg \hat{b}$$

and under convention u we obtain

$$\dot{v} \neq u \Rightarrow \dot{x} = \dot{x} + 1$$

assuming that the consequent somehow implies $\dot{v} \neq u$. Even the simple specification $\mathbf{ok} = (\dot{v} = \dot{v})$ does not respect any convention; we must instead write either $(\neg \dot{b} \Rightarrow \dot{v} = \dot{v})$ or $(\dot{v} \neq u \Rightarrow \dot{v} = \dot{v})$. All specifications become a little more complicated, with extra detail that is of no programming interest.

A stability convention is also necessary for the implementation of semi-colon: to implement $P; Q$ it must be possible to recognize termination of P 's execution in order to start Q 's execution. The usual convention used for that purpose involves part of the state (the "program counter") that programmers would prefer not to think about. No "higher-level" version of this implementation detail (such as convention u or convention b) is an improvement. Accordingly, semi-colon was defined in [3]

$$P; Q = (\forall P \Rightarrow P \circ Q)$$

with an antecedent saying only that there can be a stability convention, but not specifying any particular one. This is the strongest predicate that can be implemented as sequential execution (for proof, see [3]).

Although a stability convention is of no interest to a programmer (in fact, it is an impediment), it allows us to translate among various semantic formalisms, and compare them. That is our purpose in the remainder of this paper. In these translations and comparisons, we shall refer to the predicative semantic formalism as PS.

Dijkstra's wp

According to Dijkstra's "weakest precondition" formalism [1], a program S is defined by a predicate transformer $\text{wp}(S, R)$ that maps any postcondition R to the corresponding necessary and sufficient precondition P . Its interpretation is that, starting in a state satisfying P , execution of S will terminate in a state satisfying R .

Suppose we have a PS specification S and a postcondition R , and we want to find precondition $\text{wp}(S, R)$. First, extend the state space with at least one unstable state to provide a stability convention σ (defined to be true of exactly the unextended state space). We assume that S respects σ (if not, replace S with $(\mathbf{inv} \sigma \cdot S)$ in the following). We also assume $(\forall v \cdot R \Rightarrow \sigma)$ (if not, replace R with $R \wedge \sigma$ in the following). Now

$$\dot{\text{wp}}(S, R) = (\forall \dot{v} \cdot S \Rightarrow \dot{R}).$$

(The accent on wp signifies only that the result is a predicate in \dot{v} ; it may be deleted.) The result of this translation is a precondition P such that $(\forall v \cdot P \Rightarrow \sigma)$. The stability convention has now served its purpose; if desired, it can be thrown away by taking σ to be true, thus eliminating the unstable states.

For the reverse translation we are given $\text{wp}(S, R)$ for arbitrary R , and we

want to express S as a predicate. It is

$$S = \text{inv } \sigma \cdot \neg \dot{\text{wp}}(S, \dot{v} \neq v).$$

(For the calculation of wp , the \dot{v} are constants; the variables are unaccented. In the resulting precondition, place ` accents on the unaccented variables.) Again, we can throw away σ by letting it be true.

With the stability convention, translations between PS and wp do not lose information; they are reversible. Translation between PS and wp can also be made directly without introducing a stability convention. The translation is

$$\dot{\text{wp}}(S, R) = \nabla S \wedge (\forall \dot{v} \cdot S \Rightarrow \dot{R})$$

$$S = \neg \dot{\text{wp}}(S, \dot{v} \neq v).$$

But in one case (called **havoc** in the catalogue of semantics, later) this is not a reversible translation.

Jones's VDM

In Jones's VDM (Vienna Development Method) formalism [4], a specification is a pair (\dot{P}, R) in which \dot{P} is a precondition (predicate on the initial state) and R is a relation (predicate on the initial and final states). Its interpretation is that, starting in a state \dot{v} satisfying \dot{P} , execution will terminate in a state v such that the pair (\dot{v}, v) satisfies R .

To translate, first extend the state space with at least one unstable state to provide a stability convention σ . Then, given PS specification S respecting σ , we obtain VDM precondition \dot{P} and relation R as follows:

$$\dot{P} = \nabla S, \quad R = S.$$

In VDM, two specifications are considered equivalent if their preconditions are equivalent and their relations agree whenever \dot{v} satisfies the precondition. So we could choose any relation R such that

$$(\nabla S \wedge S) \Rightarrow R \Rightarrow (\nabla S \Rightarrow S).$$

We shall consistently choose R to be as weak as possible.

In the reverse direction, given VDM specification (\dot{P}, R) respecting σ , i.e. such that $(\forall v \cdot P \Rightarrow \sigma)$ and $(\forall \dot{v} \cdot \forall v \cdot \dot{P} \wedge R \Rightarrow \sigma)$, obtain PS specification S as

$$S = (\dot{P} \Rightarrow R).$$

Once again, with the stability convention the translations do not lose information. The translation can be made without any stability convention, but then there is the same single case in which the translation is not reversible.

Translation between VDM and wp is reversible with or without a stability

convention. Given wp ,

$$\dot{P} = \text{wp}(S, \text{true}), \quad R = \neg \text{wp}(S, v \neq v).$$

The other way, given \dot{P} and R ,

$$\text{wp}(S, Q) = \dot{P} \wedge (\forall v \cdot R \Rightarrow Q).$$

Partial Relation PR

In this formalism, a specification consists of a (possibly partial) relation R , with the following interpretation: if (\dot{v}, \dot{v}) satisfies R , then execution starting at \dot{v} must terminate, and \dot{v} is a possible final state. This formalism has been proposed by Robison [6] and independently by B. von Stengel (private communication).

An implementable PS specification S can be translated to PR by extending the state space as before, and ensuring that S is respectful. Then

$$R = \forall S \wedge S.$$

The reverse translation is

$$S = ((\exists \dot{v} \cdot R) \Rightarrow R).$$

Parnas's LD and SS

In Parnas's LD (Limited Domain) formalism [5], as in Jones's VDM, a specification is a pair (\dot{C}, R) . \dot{C} is a predicate on the initial state called the "competence", and R is a relation (predicate on the initial and final states). Its interpretation is a little different than that of VDM: if execution starts in a state \dot{v} satisfying \dot{C} , it will terminate in a state \dot{v} such that the pair (\dot{v}, \dot{v}) satisfies R ; if it starts in a state \dot{v} not satisfying \dot{C} , it will either terminate in a state \dot{v} such that the pair (\dot{v}, \dot{v}) satisfies R or it will fail to terminate. Unlike VDM, specifications that differ only when \dot{v} lies outside the competence set are not considered equivalent.

The translation between PS and LD is similar to the translation between PS and VDM for specifications that respect a stability convention. LD differs from VDM in its ability to be disrespectful: an LD specification can say that, for some initial state, nontermination and termination in a limited set of final states are acceptable, but termination in any other final state is unacceptable.

$$\dot{C} = (\forall \dot{v} \cdot S \Rightarrow \dot{\sigma}), \quad R = S \wedge \dot{\sigma}$$

$$S = (\dot{C} \vee \dot{\sigma} \Rightarrow R).$$

Parnas has also considered a relational semantics that he has called "standard semantics" SS. It is similar to LD but the competence is determinable from the relation as

$$\dot{C} = \exists \dot{v} \cdot R$$

and so we can dispense with it. SS can be seen as a variation of PR. In PR, if for \dot{v} there is no \dot{v} such that (\dot{v}, \dot{v}) satisfies R , then execution starting at \dot{v} is arbitrary; in SS, nontermination is mandatory.

$$\begin{aligned}
 R &= \text{if } \forall \dot{v} \cdot (\forall \dot{v} \cdot S \Rightarrow \dot{\sigma}) \vee (\forall \dot{v} \cdot S \Rightarrow \neg \dot{\sigma}) \\
 &\quad \text{then } S \wedge \dot{\sigma} \\
 &\quad \text{else inexpressible} \\
 S &= (\dot{\sigma} \Rightarrow (\exists \dot{v} \cdot R)) \wedge ((\exists \dot{v} \cdot R) \Rightarrow R).
 \end{aligned}$$

Catalogue of Semantics

We now present a dozen specifications and connectives, each in the six formalisms we have just described. Those that are basic (not composed from others) are shown both with and without a stability convention. The reader should compare them for simplicity and ease of expression. Our comparative comments follow.

specification called **abort**, **chaos** or **disaster**: arbitrary behavior

	with σ	without σ
PS	true	true
wp(S, R)	false	false
VDM	false, true	false, true
PR	false	false
LD	false, true	false, true
SS	inexpressible	inexpressible

specification called **havoc** or **chance**: terminating but otherwise arbitrary behavior

	with σ	without σ
PS	$\dot{\sigma} \Rightarrow \dot{\sigma}$	inexpressible
wp(S, R)	$\forall \dot{v} \cdot \sigma \Rightarrow R$	$\forall \dot{v} \cdot R$
VDM	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{\sigma}$	true, true
PR	$\dot{\sigma} \wedge \dot{\sigma}$	true
LD	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{\sigma}$	true, true
SS	inexpressible	true

specification called **skip**, **ok** or **continue**

	with σ	without σ
PS	$\dot{\sigma} \Rightarrow \dot{v} = \dot{v}$	$\dot{v} = \dot{v}$
wp(S, R)	$\sigma \wedge R$	R
VDM	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{v} = \dot{v}$	true, $\dot{v} = \dot{v}$
PR	$\dot{\sigma} \wedge \dot{v} = \dot{v}$	$\dot{v} = \dot{v}$
LD	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{v} = \dot{v}$	true, $\dot{v} = \dot{v}$
SS	inexpressible	$\dot{v} = \dot{v}$

specification called **miracle**

	with σ	without σ
PS	$\neg \dot{\sigma}$	false
wp(S, R)	σ	true
VDM	$\dot{\sigma}, \neg \dot{\sigma}$	true, false
PR	inexpressible	inexpressible
LD	$\dot{\sigma}, \neg \dot{\sigma}$	true, false
SS	inexpressible	inexpressible

assignment $x := e$ in two variables x, y , assuming e is evaluable

	with σ	without σ
PS	$\dot{\sigma} \Rightarrow \dot{x} = \dot{e} \wedge \dot{y} = \dot{y} \wedge \dot{\sigma}$	$\dot{x} = \dot{e} \wedge \dot{y} = \dot{y}$
wp(S, R)	$\sigma \wedge R[x: e]$	$R[x: e]$
VDM	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{x} = \dot{e} \wedge \dot{y} = \dot{y} \wedge \dot{\sigma}$	true , $\dot{x} = \dot{e} \wedge \dot{y} = \dot{y}$
PR	$\dot{\sigma} \wedge \dot{x} = \dot{e} \wedge \dot{y} = \dot{y} \wedge \dot{\sigma}$	$\dot{x} = \dot{e} \wedge \dot{y} = \dot{y}$
LD	$\dot{\sigma}, \dot{\sigma} \Rightarrow \dot{x} = \dot{e} \wedge \dot{y} = \dot{y} \wedge \dot{\sigma}$	true , $\dot{x} = \dot{e} \wedge \dot{y} = \dot{y}$
SS	inexpressible	$\dot{x} = \dot{e} \wedge \dot{y} = \dot{y}$

sequential composition: satisfy first P then Q

PS	$\nabla P \Rightarrow P \circ Q$
wp(S, R)	wp(P, R) \wedge wp(Q, R)
VDM	$\dot{P}_P \wedge (\forall \dot{v} \cdot R_P \Rightarrow \dot{P}_Q), R_P \circ R_Q$
PR	$(\forall \dot{v} \cdot P(\dot{v}, v) \Rightarrow \exists \dot{v} \cdot Q(\dot{v}, v)) \wedge P \circ Q$
LD	$\dot{C}_P \wedge (\forall \dot{v} \cdot R_P \Rightarrow \dot{C}_Q), R_P \circ R_Q$
SS	if $\forall \dot{v} \cdot (\exists \dot{v} \cdot P \circ Q) \Rightarrow (\forall \dot{v} \cdot P(\dot{v}, v) \Rightarrow \exists \dot{v} \cdot Q(\dot{v}, v))$ then $P \circ Q$ else inexpressible

deterministic choice: if b is initially true then satisfy P else satisfy Q

PS	$\dot{b} \wedge P \vee \neg \dot{b} \wedge Q$
wp(S, R)	$\dot{b} \wedge \text{wp}(P, R) \vee \neg \dot{b} \wedge \text{wp}(Q, R)$
VDM	$\dot{b} \wedge \dot{P}_P \vee \neg \dot{b} \wedge \dot{P}_Q, \dot{b} \wedge R_P \vee \neg \dot{b} \wedge R_Q$
PR	$\dot{b} \wedge P \vee \neg \dot{b} \wedge Q$
LD	$\dot{b} \wedge \dot{C}_P \vee \neg \dot{b} \wedge \dot{C}_Q, \dot{b} \wedge R_P \vee \neg \dot{b} \wedge R_Q$
SS	$\dot{b} \wedge P \vee \neg \dot{b} \wedge Q$

nondeterministic choice: satisfy either P or Q

PS	$P \vee Q$
wp(S, R)	wp(P, R) \wedge wp(Q, R)
VDM	$\dot{P}_P \wedge \dot{P}_Q, \dot{P}_P \wedge \dot{P}_Q \Rightarrow R_P \vee R_Q$
PR	$(\exists \dot{v} \cdot P) \wedge (\exists \dot{v} \cdot Q) \wedge (P \vee Q)$
LD	$\dot{C}_P \wedge \dot{C}_Q, R_P \vee R_Q$
SS	if $\forall \dot{v} \cdot (\exists \dot{v} \cdot P) = (\exists \dot{v} \cdot Q)$ then $P \vee Q$ else inexpressible

joint satisfaction: satisfy both P and Q

PS	$P \wedge Q$
wp(S, R)	$\text{wp}(P, R) \vee \text{wp}(Q, R)$
VDM	$\hat{P}_P \vee \hat{P}_Q, R_P \wedge R_Q$
PR	if $\forall \dot{v} \cdot (\exists \dot{v} \cdot P \wedge Q) \vee (\neg \exists \dot{v} \cdot P) \vee (\neg \exists \dot{v} \cdot Q)$ then $((\exists \dot{v} \cdot P \wedge Q) \Rightarrow P \wedge Q)$ $\wedge ((\neg \exists \dot{v} \cdot P) \vee (\neg \exists \dot{v} \cdot Q) \Rightarrow P \vee Q)$ else inexpressible
LD	$\hat{C}_P \vee \hat{C}_Q, R_P \wedge R_Q$
SS	if $\forall \dot{v} \cdot (\exists \dot{v} \cdot P \wedge Q) \vee (\neg \exists \dot{v} \cdot P) \wedge (\neg \exists \dot{v} \cdot Q)$ then $P \wedge Q$ else inexpressible

ordering: P satisfied if Q satisfied

PS	$P \Leftarrow Q$
wp(S, R)	$\forall R \cdot \text{wp}(P, R) \Rightarrow \text{wp}(Q, R)$
VDM	$\hat{P}_P \Rightarrow \hat{P}_Q \wedge (R_P \Leftarrow R_Q)$
PR	$(\exists \dot{v} \cdot P) \Rightarrow (\exists \dot{v} \cdot Q) \wedge (P \Leftarrow Q)$
LD	$\hat{C}_P \Rightarrow \hat{C}_Q \wedge (R_P \Leftarrow R_Q)$
SS	$((\exists \dot{v} \cdot P) = (\exists \dot{v} \cdot Q)) \wedge (P \Leftarrow Q)$

variable declaration: introduce fresh local variable x into P

PS	$\exists \dot{x} \cdot \exists \dot{x} \cdot P$
wp(S, R)	$\forall x \cdot \text{wp}(P, R)$
VDM	$\forall \dot{x} \cdot \hat{P}_P, \exists \dot{x} \cdot \exists \dot{x} \cdot R_P$
PR	$\exists \dot{x} \cdot \exists \dot{x} \cdot P$
LD	$\forall \dot{x} \cdot \hat{C}_P, \exists \dot{x} \cdot \exists \dot{x} \cdot R_P$
SS	if $\forall \dot{v} \cdot \forall \dot{v} \cdot (\exists \dot{x} \cdot \exists \dot{x} \cdot P) \Rightarrow (\forall \dot{x} \cdot \exists \dot{x} \cdot P)$ then $\exists \dot{x} \cdot \exists \dot{x} \cdot P$ else inexpressible

recursion: to satisfy **loop** $P: S(P)$, satisfy $S(\text{loop } P: S(P))$

PS	$\forall i \cdot S^i(\text{true})$
wp(S, R)	$\exists i \cdot \text{wp}(S^i(\text{abort}), R)$
VDM	$\exists i \cdot \hat{P}_{S^i(\text{abort})}, \forall i \cdot R_{S^i(\text{abort})}$
PR	not easily expressible
LD	not easily expressible
SS	inexpressible.

Comparative Comments

It is questionable whether a semantic formalism that talks only about the initial input state and the final output state, without any interactive input or output during the computation, can be considered adequate. But all the formalisms considered here are of that sort, so it is not a point of comparison.

All of the formalisms provide a way of saying, for each input state, "specification P requires termination". They are:

PS	∇P
wp	$\text{wp}(P, \text{true})$
VDM	\hat{P}_P
PR	$\exists \dot{v} \cdot P$
LD	\hat{C}_P
SS	$\exists \dot{v} \cdot P.$

The negation means, of course, "specification P does not require termination". Only two of the formalisms provide a way of saying "specification P requires nontermination". They are:

LD	$\neg \hat{C}_P \wedge \neg \exists \dot{v} \cdot P$
SS	$\neg \exists \dot{v} \cdot P.$

In PS, ∇P says that P is determined; to know that a computation satisfies a determined specification, one must observe its final state (so it must have one). But to know that a computation satisfies an undetermined specification, i.e. when $\forall \dot{v} \cdot P$, it is not necessary to observe its final state (so it does not need to have one).

One of the formalisms, wp, is not based on a relation. This makes it awkward to state that the output is to be related to the input. To say that x is to be increased by 1, we may write something like

$$\text{wp}(S, x = X + 1)$$

accompanied by words that say X is the initial value of x . But formally this allows the unintended solution

$$X := 3; x := 4$$

for S . To say what is intended, we must be able to distinguish, in the formalism, between initial and final values.

VDM and PS are very close, as seen by the simplicity of the translation. In practice, it often makes sense to think of the precondition separately from the relation, as in VDM. One writes them down with a comma between, as in VDM, or with an implication sign, as in PS. VDM offers us a freedom in stating the relation, and we chose to make it as weak as possible. Consequently,

$$\forall \dot{v} \cdot \nabla R \Rightarrow \hat{P}.$$

This means that the precondition is entirely superfluous whenever ∇R is true. It gives us further information only for those inputs where ∇R is false. If, for some input, all outputs are acceptable, the precondition tells us whether termination is required. Its only service is to distinguish, for each input state, **chaos** from **havoc**. In judging the value of this service against its cost, keep in mind the fact that termination, by itself, is not observable.

LD differs from VDM in its ability to be disrespectful: an LD specification can say that, for some initial state, nontermination and termination in a limited set of final states are acceptable, but termination in any other final state is unacceptable. Such a specification is very easy to implement: ignore the final states and deliver a nonterminating loop. Hence the ability to be disrespectful is of no practical value in programming. PS, wp, VDM, and PR all hold that (without communications or intermediate results) infinite loops are useless, and there is no point in being particular about possible final states while allowing an infinite loop. One may take the position that LD can describe more mechanisms than PS, wp, VDM, or PR: it can describe those that can nondeterministically either terminate in a particular final state or not terminate. Or, one may instead take the position that mechanisms are deterministic, and that nondeterminism is a property of specifications. In that case, LD is no more descriptive than the other formalisms.

To a specifier, the most important connective is probably joint satisfaction. It is very common to specify by parts, stating properties that must be jointly satisfied. A specifier will therefore appreciate the formalism that makes joint satisfaction simplest.

To a programmer, the most important connective is undoubtedly the ordering, because that is the criterion for correctness of a program. Indeed, the entire programming process is to transform a specification, little by little, to a program, and at each stage to ensure that the new is properly related to the old by the ordering. A programmer will therefore appreciate the formalism that makes the ordering simplest.

Of all these formalisms, PR fits the class of computations best – in fact, perfectly. All implementable specifications are expressible, and all unimplementable specifications are not. That was its design criterion. To achieve this, PR significantly complicates the expression of its connectives, particularly joint satisfaction.

Alone among the formalisms, PS can be criticized for the fact that its semi-colon is not always associative. The counter-example creates, using semi-colon, what would be **havoc** in another formalism. For example, in one integer variable x ,

$$\begin{aligned} (x \neq x; x \neq x); x = 0 &= \text{true} \\ x \neq x; (x \neq x; x = 0) &= x = 0. \end{aligned}$$

Our intuition that semi-colon should be associative is simply this: executing first P and then $(Q; R)$ should be the same as executing first $(P; Q)$ and then R . This makes sense when execution of P and $(P; Q)$ terminate. But if P 's execution fails to terminate, what does "and then $(Q; R)$ " mean? If $(P; Q)$'s execution fails to terminate, what does "and then R " mean? It is a theorem of PS that semi-colon is associative under all circumstances called for by intuition:

$$\forall P \wedge \forall (P; Q) \Rightarrow (P; (Q; R) = (P; Q); R).$$

In fact, the stronger theorem

$$\forall (P; Q) \Rightarrow \forall P \wedge (P; (Q; R) = (P; Q); R)$$

holds.

More Power

We now present one more semantic formalism, called PF (Powerful Formalism). As in five of the formalisms presented so far, the heart of PF is a relation between input \hat{v} and output \hat{v} . Like wp, VDM and LD, it has further information to impart. In PF a specification is a sextuple $(\hat{w}g, \hat{w}lg, \hat{w}p, \hat{w}lp, R, \hat{s}t)$ in which the first four components are predicates on the initial state, R is the relation, and $\hat{s}t$ is a predicate on the final state. These components are to be interpreted as follows.

$\hat{w}g$: true of an initial state from which the computation is required to progress, false of an initial state in which the computation is allowed to get stuck.

$\hat{w}lg$: true of an initial state from which the computation is allowed to progress, false of an initial state in which the computation is required to get stuck.

$\hat{w}p$: true of an initial state for which the computation is required to terminate, false of an initial state for which the computation is allowed to run forever.

$\hat{w}lp$: true of an initial state for which the computation is allowed to terminate, false of an initial state for which the computation is required to run forever.

R : this relation describes the allowed transitions from initial to final state.

$\hat{s}t$: true of a final state iff termination is successful.

Here is a small sample of basic specifications in PF

```

chaos  = false, true,  false, true,  true,  false
havoc  = true,  true,  true,  true,  true,  false
random = true,  true,  true,  true,  true,  true
break  = false, false, false, false, false, false
miracle = true,  true,  true,  true,  false, true
magic  = false, true,  false, true,  false, true
ok      = true,  true,  true,  true,   $\hat{v} = \hat{v}$ , true
skip   = false, true,  true,  true,   $\hat{v} = \hat{v}$ , true
spin    = true,  true,  false, false,  $\hat{v} = \hat{v}$ , false

```

In this formalism, we are able to express not only **chaos** and **havoc**, but **random** and **break** as well, and we can distinguish **ok** from **skip**. An advocate of PF might criticize all weaker formalisms for being unable to make these distinctions. One might well ask if **random** and **break** are useful for any purpose; the answer is that they are just as useful as **havoc**. One might be forgiven for wondering how we are to observe “progress” and “success”; the answer is similar to that for “termination”.

Is PF powerful enough? We gave it components that sound almost reasonable; properties like them have all been seriously discussed in the literature. We were tempted to give it a few more components, such as “strangeness” and “charm”, to make our point more obvious. If we wish, each of these properties can be observed indirectly via a convention, which is an interpretation of those things that are directly observable. But what is their purpose?

We have kept PS simple. Those who are uncomfortable without an explicit termination indicator are invited to write their programs and other specifications using their favorite termination convention. Then the initial states for which the computation is required to terminate are those for which $(\forall v \cdot S \Rightarrow \phi)$, i.e. for which the specification implies termination. PS provides the ability to express any convention, but it has none built in.

Conclusion

In all practical circumstances, we can and do observe termination: by means of a termination convention. This is so standard that it may be difficult to realize we are observing a convention, and not termination itself.

We have said that "termination" does not have a physical meaning. We now point out that it does not always have a logical meaning, either. For any semantic formalism and any proof theory, there is a program P such that " P terminates" can neither be proved, nor disproved. There is a model in which it is true, and another in which it is false. For each such program P , the proof theory can be enriched and the model theory constrained to make " P terminates" mean whichever we prefer: true or false. To those who have it firmly fixed in their heads that "termination" means something, this will seem paradoxical.

The simplest of the formalisms, having the properties most convenient for both specifiers and programmers, is PS without a stability convention. Its one "limitation" is its inability to distinguish between **chaos** and **havoc**. But no physical experiment can distinguish between them, either. The distinction is a creation of the more complex formalisms.

Acknowledgement. We thank Cliff Jones and an anonymous referee for helpful suggestions.

References

1. E.W. Dijkstra: A Discipline of Programming. Englewood Cliffs, N.J.: Prentice-Hall 1976
2. E.C.R. Hehner: Predicative Programming. *Commun. ACM* **27**, 134–151 (1984)
3. E.C.R. Hehner, L.E. Gupta, A.J. Malton: Predicative Methodology. *Acta Inf.* **23**, 487–505 (1986)
4. C.B. Jones: Systematic Software Development using VDM. London: Prentice-Hall Int. 1986
5. D.L. Parnas: A generalized control structure and its formal definition. *C. ACM* **26**, 572–581 (1983)
6. W.A. Robison: A Deterministic Trace Semantics for Communicating Sequential Processes. M.A.Sc. Thesis, University of Toronto, 1986

Received October 2, 1986/July 7, 1987