





# Challenges for Formal Semantic Description: Responses from the Main Approaches

Cliff B. Jones<sup>(✉)</sup>  and Troy K. Astarte 

School of Computing, Newcastle University, Newcastle upon Tyne NE4 5TG, UK  
`cliff.jones@ncl.ac.uk`

**Abstract.** Although there are thousands of programming languages, most of them share common features. This paper reviews some key underlying language concepts and the challenges they present to the task of formally describing language semantics. The responses to these challenges in operational, axiomatic and denotational approaches to semantic description are reviewed. There are interesting overlaps between these responses; similarities are exposed even where accidental notational conventions disguise them so that essential differences can be pinpointed. Depending on the objectives of writing a formal semantic description of a language, one or other approach might be considered the best choice. An argument is made for increasing the use of formal semantics in language design and here it is suggested that the operational approach is the most viable for a complete language description.

## 1 Introduction

There are a number of different approaches to recording formal descriptions of the semantics of programming languages, but most can be placed into one of three styles: operational, denotational, or axiomatic. Any approach to describing semantics formally must find ways to tackle a set of challenges derived from common features in programming languages, such as nested blocks or concurrency. In this paper, an initially simple illustrative language is described using all three approaches and remarks are made about how they address the particular challenges. It is interesting to note the degrees of similarity present given the apparent conceptual differences between approaches.

The paper begins by setting out some reasons for considering semantics and introducing the kernel of the example language. Simple applicative languages are considered first and some conclusions are drawn that are relevant to imperative languages. Throughout the paper, new features for the example language are considered and the formal semantic descriptions of these features are discussed. Finally, a concurrent, object-oriented language is introduced as a vehicle to illustrate the combination of all the features covered; an operational semantics for such a language is outlined.

This is not primarily intended to be a historical paper; readers interested in such a view of formal semantics could read [AJ18] which examines four early semantic descriptions of ALGOL 60 and draws some conclusions. A more complete treatment of the history of programming language semantics is presented in Astarte’s PhD thesis [Ast19]. Nor is the objective here to provide a tutorial on semantic description formalisms; the reader is assumed to have some previous contact with the subject. The aim here is to look beyond the trivial language features that are easily handled by any formal approach.

## 1.1 Why Describe Semantics Formally

It is worth beginning by reviewing the reasons for describing the semantics of programming languages. Unlike natural languages, programming languages are formal objects which means they follow a fixed (and relatively small) set of rules that govern their structure and behaviour.

It is essential that the different users of a language, from programmers through standard writers to compiler creators, all share a precise understanding of these rules.<sup>1</sup> Natural language can be (and is) used for this purpose, but words are always ambiguous and can all too easily lead to contradictions or omissions. Therefore, formality is frequently utilised—and even in natural language descriptions, the careful wording required ultimately results in formality of approach regardless of notation [Tur09]. Another advantage to the use of formalism is that it can help ensure completeness: if there is a form to be followed for every language construct, the chances of accidentally omitting part of a language is significantly lowered.

This is not to suggest that a formal description always defines one unique result for a program in a language: it is often necessary to leave certain parts of the description undefined in order to allow for implementation specifics and non-determinism at run time. Carefully delineating these areas of non-definition is, however, essential.

In addition to being formal, a useful programming language semantics must also be *tractable*—it must enable proofs to be made about the language itself, about the correctness of implementations of the language and about programs written in the language [Bur66]. Ideally, a good semantics allows the proof of deep properties, some of which are relied upon in compiler optimisation. Different approaches to semantics tend to make different properties easier to prove than others [Gor75].

Arguably an even more important use of formal semantics is in the design of programming languages: there exist thousands of programming languages, most of which are sadly lamentable;<sup>2</sup> even the best often exhibit *feature interaction*

<sup>1</sup> See, for example, the work of the IBM Laboratory Vienna on producing formal definitions of PL/I for use in compiler writing, such as [BBH+74, Jon76].

<sup>2</sup> In the paper ‘Hints of Programming Language Design’, Tony Hoare had the following to say in conclusion: “This paper has given many practical hints on how *not* to design a programming language. It has even suggested that many recent languages have followed these hints” [Hoa73].

where features that are useful and straightforward when taken separately lead to incomprehensible behaviour when combined. The use of a formal semantics during the creation of a language—ideally, before even any syntax is created—can contribute greatly to the simplicity and clarity of the resultant language. Unfortunately, formal semantics has typically been applied *post facto* to extant languages.<sup>3</sup> Arguments for the use of semantics in the design of languages are given in Sect. 7.

The choice of semantic description approach is often motivated by the intended use of the semantics. Received wisdom generally holds that operational semantics is most useful to compiler writers, denotational to the language designer and axiomatic semantics in program verification. However, some writers have pointed out that the distinction is not always as clear cut as this [Ame89].

Of course, the challenge of describing the semantics of a modern programming language is far greater than for, say, first-order predicate calculus. Researchers have learnt what they can from previous work by logicians and carried these lessons forward: the extensions involved are challenging and interesting.

## 1.2 Main Approaches

The main focus in this paper is on operational, axiomatic and denotational semantics; Sect. 2 illustrates the differences in these approaches on a core language but it is worth briefly characterising the approaches here.

An operational semantics describes the meaning of a language in terms of an *abstract interpreter* that takes a program and a starting state and computes allowed final states. Typically, the interpreter will be defined in terms of subfunctions for each construct in the language. Ideally, the *states* of the interpreter should be chosen to eschew unnecessary details.

The essence of a denotational semantics is to map a language into some space of mathematically tractable objects. For simple programming languages these objects are mathematical functions from states to states. Denotational descriptions present a series of mappings from program constructs into these functions. A key feature is the notion that the mapping should be *homomorphic*: the function denoted by a program segment should be composed from the denotations of its components.

The previous two approaches both make the notion of state explicit and can thus be categorised as *model-based*. In contrast, property-oriented descriptions attempt to fix semantics without an explicit state.<sup>4</sup> An axiomatic semantics contains axioms and rules of inference that define a set of *judgements*. In Floyd-Hoare semantics of procedural languages, the judgements are triples in which

<sup>3</sup> Encouraging exceptions include the Turing programming language [HMRC87], Standard ML [HMT87], and SPARK-Ada [CG90]. Furthermore, formal semantics played an important role in the development of full Ada [BO80]. Formal description was also utilised in the standards for Modula-2 [Woo93] and PL/I [ANS76].

<sup>4</sup> *Algebraic semantics* can also be viewed as property-oriented and is briefly discussed in Sect. 7.1.

the middle component is a text in the language being described; the first and third components are predicates. The interpretation of such a triple is that if the first predicate (the pre condition) is satisfied and the text is executed to termination, then after execution the post condition will be true.

Here the notion of state is only implicit in the meta-variables used within the assertions. Axiomatic semantics is particularly concerned with proving properties of programs and, if an axiomatic specification of a language allows the proving of any true property (and no false property) of a program construct, then the construct is considered fully specified [Pag81]. If every part of the language is specified in this way, then the specification constitutes a semantics of the language. In practice, it turns out to be difficult to fully specify large-scale programming languages purely by axioms.

### 1.3 Applicative Languages

The majority of this paper is concerned with *imperative* programming languages (as characterised in Sect. 2). There are, however, some interesting semantic description techniques that can be carried over from handling *applicative* languages. Two common challenges are that the languages whose semantics are to be given have an unbounded number of admissible texts and that comprehensibility of the semantic description is a major objective.

One class of applicative programming language is functional programming languages and these – at least if they are *purely* functional – avoid some of the challenges that have to be faced with the semantics of languages that feature assignment-like constructs. Assignments require some model of storage, usually considered as an abstract meta-notion *state*; avoiding assignment allows programs in functional languages to be reasoned about as though they are conventional recursive functions. There might, of course, be a performance penalty in using purely functional languages, but that discussion is beyond the scope of the present paper.

It is important to remember that all programming languages provide a repertoire of basic operators and, crucially, put in the hands of programmers ways to express functions that extend this repertoire. Thus a programmer might write a program that computes factorial using only basic arithmetic operators; more ambitiously, a program for inverting matrices can be written in a language that has no such operator.

A first-order predicate language is a simple and traditional applicative language and discussing how its semantics can be recorded facilitates deriving messages that are taken forward to the subsequent sections of this paper. Starting with purely propositional expressions, a semantic function could be written that recurses through the structure of the expressions,<sup>5</sup> building up the meaning of the expression as a whole by combining the meaning of its parts. Ultimately, this function must rely on an association of the propositional identifiers with

---

<sup>5</sup> This task would be made easier with the use of an abstract syntax, a concept discussed later in this paper.

truth values. As with predicate calculus, there must be a way to determine the meaning of any predicates or functions. It is important to observe that these two sorts of associations remain fixed and can be stored in some form of static *environment*.

There are, of course, other ways of tackling the semantics of logical languages. In an equivalence-based strategy, some operators can be defined in terms of others (e.g.  $p \Rightarrow q$  can be defined as  $\neg p \vee q$ ). No matter the strategy used, there must still be a minimum set of basic operators (e.g. the Sheffer stroke).

Classical axiomatisations (such as that in [Men64, Sect. 1.4]) are unintuitive but natural deduction rules like those presented in [Pra65] provide both a semantics and some intuition as to how to reason about expressions in logical languages.

The responses to be carried forward to the review of semantic description techniques for imperative languages are then:

- Environments—what information is stored about identifiers; in what form; and how distinction is made between different denotations e.g. identifier-value and function-definition pairs.
- Fundamental bases of meaning—saying one has, for example, a *Boolean Algebra* doesn't fix (all of) the semantics because multiple models of such algebras exist.
- Understandability of description—as with deduction systems, semantic descriptions should be evaluated for intuition and usability for reasoning.

## 1.4 A Core Imperative Language

A basic challenge to be faced, even before addressing the semantics of a language, is to delimit the admissible utterances of the language. Although normally presented in two dimensional layout, it is still common to think of programs as strings of characters. Some version of Backus-Naur Form notation is adequate to define the set of (context-free) strings of most programming languages: this is known as *concrete syntax*. However, following Christopher Strachey's advice to "know what you need to say before deciding how to write it down", semantic descriptions can be based instead on an *abstract syntax*. This approach follows John McCarthy's proposal [McC63] although VDM notation is employed below.<sup>6</sup> The advantages of using an abstract syntax over concrete may be less apparent for a small language like the one considered here but for large languages, especially those with multiple syntactic forms of the same semantic construct, the benefits become more apparent. Use of abstract syntax shows concern with the *structure* of the language (rather than its form). The higher level of abstraction meshes nicely with more abstract semantic approaches; however, following tradition, examples of axiomatic semantics below are built around concrete syntax.

---

<sup>6</sup> The use of VDM notation should present the reader with no difficulty: it has been widely used for decades and is the subject of an ISO standard; one useful reference is [Jon90].

Figure 1 contains the abstract syntax of the simple core of the language discussed in this paper. Later sections in the paper add to this core to illustrate more complex language concepts and the challenges inherent in modelling these features.

$$\begin{aligned}
 \text{Program} &:: \text{types} : \text{Id} \xrightarrow{m} \text{ScalarType} \\
 &\quad \text{body} : \text{Stmt} \\
 \text{ScalarType} &= \text{INT} \mid \text{BOOL} \\
 \text{Stmt} &= \text{Assign} \mid \text{If} \mid \text{While} \mid \text{Compound} \mid \dots \\
 \text{Assign} &:: \text{lhs} : \text{Id} \\
 &\quad \text{rhs} : \text{Expr} \\
 \text{If} &:: \text{test} : \text{Expr} \\
 &\quad \text{then} : \text{Stmt} \\
 &\quad \text{else} : \text{Stmt} \\
 \text{While} &:: \text{test} : \text{Expr} \\
 &\quad \text{body} : \text{Stmt} \\
 \text{Compound} &:: \text{Stmt}^*
 \end{aligned}$$

**Fig. 1.** Abstract syntax of a core language

Even before getting to the semantic approaches *per se*, it is worth noting that there are differences as to how context-dependent checks (e.g. required consistency between uses and declarations of names) are recorded. These can be handled within semantics (i.e. dynamically), but it is normally more fruitful to handle these issues statically. Such static checks are called *context conditions* after van Wijngaarden *et al.* in the ALGOL 68 Report [vWMPK69]. Various methods for defining these kinds of checks have been developed by van Wijngaarden (two-level grammars), Knuth (attribute grammars [Knu68]), and researchers at the IBM Hursley Laboratory (dynamic syntax [HJ73]); a more thorough study would include [GP99] or [Pie02] on *type theory*. Full exploration of this topic is beyond the scope of the current paper.

Context conditions in the VDM style are generally written as predicates that determine whether an object of the abstract syntax is *well-formed* with respect to some type information. For the language whose abstract syntax is given in Fig. 1, these predicates would have the signature  $wf\text{-Construct} : \text{Construct} \times \text{TypeMap} \rightarrow \mathbb{B}$  and use an abstract *TypeMap* object of the type  $\text{Id} \xrightarrow{m} \text{ScalarType}$  (a finite, constructed, function) that maps identifiers to their types.<sup>7</sup> In this simple case, the *TypeMap* is a direct copy of that in the *Program*. These functions generally check that the types assigned to variables match the variable declaration and

<sup>7</sup> The use of the type name *ScalarType* prepares the way for modelling compound types such as arrays below.

that inappropriate types are not used in expressions (for example, in an *If* statement, the *test* part must be of type `BOOL`). For constructs that contain sub-components, each such component must also be well formed.

## 2 Imperative (Deterministic) Languages

The identifying feature of an imperative programming language is that it provides statements that change things. What is affected differs between languages: changes might be updating a database or moving the position of part of a robot. Here the discussion focusses on the challenge of modelling assignments to variables but the same principles apply to other kinds of command as long as a suitable abstract model can be created for the target of the changes.

Assignments to variables destroy *referential transparency*: the value associated with an identifier changes during execution; values previous to an assignment are destroyed. Furthermore, the order in which statements are executed becomes important. An imperative program achieves its effect by executing a sequence of assignments; language features such as conditionals and loops merely orchestrate their execution.

As in applicative languages, programs make it possible to compute results that are not directly available as operators of the language. It therefore follows that a subsidiary challenge is to provide tractable ways of reasoning about the meaning of imperative programs whose specifications include operators that are not basic to the language and which achieve their effect using destructive assignments.

### 2.1 Operational Approach

John McCarthy was one of the first to present an operational approach to defining the semantics of programming languages. In his definition of ‘Micro-ALGOL’ [McC66], he described the approach as “defining a function ... giving the state ... that results from applying the program ... to the [initial state]”. McCarthy was also careful to point out in his earlier paper on the topic [McC63] that this is an *abstract function*, because the language in which it is expressed is more abstract than either the language being described or, say, machine assembler code. This approach to semantics is now commonly referred to as an *abstract interpreter* because it interprets the various constructs of the language under discussion.

The core idea of operational semantics remains the same as when McCarthy first proposed it: meaning is given to a language with an abstract interpreter defined in terms of changes to abstract states. The capital Greek letter  $\Sigma$  is commonly used for the set of such states and, in simple cases, particular states directly associate identifiers with values such as Booleans or integers:

$$\Sigma = Id \xrightarrow{m} \text{ScalarValue}$$

$$\text{ScalarValue} = \mathbb{B} | \mathbb{Z}$$

As observed above, the key property of an imperative language is that assignments can change the state. An interpretation function for statements would take as parameters an (abstract) program and a state; its result is a final state. Historically, McCarthy [McC66] and even the early Vienna operational descriptions (such as the VDL descriptions of PL/I [Lab66]) did write such interpretation functions. In the current paper, the *Structural Operational Semantics* (SOS) style of [Plo81] is used uniformly since this notation copes with non-determinism (cf. Section 4.1) and can thus be used for all of the operational descriptions discussed.

SOS rules like the one below for assignment can be read like a classic interpretation function, when considered in a clockwise manner from bottom left, and this often feels more natural when looking at deterministic languages. However, it is important to remember that SOS rules are in fact *inference rules*: above the line is a series of premises which must all be true for the rule to apply; below the line is the conclusion. Each rule indicates a relation between the state before computation and the state afterwards, given that a series of conditions holds; it records a way of judging whether a particular computation is valid. This distinction becomes important when considering non-deterministic languages, as in Sect. 4.1.

The basic judgements are relations (thus their signatures use powersets) between pairs of program text and pre-state, and post-computation state. The relation for *statements* is:

$$\xrightarrow{st}: \mathcal{P}((Stmt \times \Sigma) \times \Sigma)$$

The precise way in which each statement in a program is interpreted obviously depends on the type of statement so one way to present a description would be to write an interpretation function by cases. However, it is more convenient to use pattern matching<sup>8</sup> and this approach is used in both operational semantics and denotational semantics:

$$\frac{(rhs, \sigma) \xrightarrow{ex} v}{(mk-Assign(lhs, rhs), \sigma) \xrightarrow{st} \sigma \uparrow \{lhs \mapsto v\}}$$

(The judgements for *expression* evaluation ( $\xrightarrow{ex}$ ) are described below.)

Conditional statements are interpreted by cases as follows:

$$\frac{\begin{array}{l} (test, \sigma) \xrightarrow{ex} \mathbf{true} \\ (then, \sigma) \xrightarrow{st} \sigma' \end{array}}{(mk-If(test, then, else), \sigma) \xrightarrow{st} \sigma'} \quad \frac{\begin{array}{l} (test, \sigma) \xrightarrow{ex} \mathbf{false} \\ (else, \sigma) \xrightarrow{st} \sigma' \end{array}}{(mk-If(test, then, else), \sigma) \xrightarrow{st} \sigma'}$$

Interpreting iterative statements is slightly more involved:

$$\frac{\begin{array}{l} (test, \sigma) \xrightarrow{ex} \mathbf{true} \\ (body, \sigma) \xrightarrow{st} \sigma' \end{array}}{(mk-While(test, body), \sigma') \xrightarrow{st} \sigma''} \quad \frac{\begin{array}{l} (test, \sigma) \xrightarrow{ex} \mathbf{false} \end{array}}{(mk-While(test, body), \sigma) \xrightarrow{st} \sigma}$$

<sup>8</sup> Each VDM record type has an associated constructor function equivalent to a type constructed by this function—thus  $mk-Assign : Id \times Expr \rightarrow Assign$  can be used to distinguish the appropriate subset of *Stmt* in a pattern matching context.



Notice that the state used in the third premise is the one produced from an interpretation of the body—thus a convergence towards termination may occur. The issue of non-terminating loops is addressed below.

The basic notion of state used above plays the same role as the environment in a functional language and an evaluation function can be defined to determine the values of expressions.

$$eval : Expr \times \Sigma \rightarrow ScalarValue$$

The *eval* function above can be rewritten as a relation:<sup>9</sup>

$$\xrightarrow{ex} : \mathcal{P}((Expr \times \Sigma) \times ScVal)$$

which can be split by the cases in its syntactic classes

$$\frac{e \in Id}{(e, \sigma) \xrightarrow{ex} \sigma(e)}$$

$$\frac{(e1, \sigma) \xrightarrow{ex} v1 \quad (e2, \sigma) \xrightarrow{ex} v2}{(mk-Expr(e1, PLUS, e2), \sigma) \xrightarrow{ex} v1 + v2}$$

Other cases should be obvious.

This seemingly simple description actually fixes an important property of the language: the process of evaluating an expression is shown not to change the state (i.e. the values of variables)—the same  $\sigma$  is used throughout. Although the key feature of functions is not addressed until Sect. 3, it is important to note that functions with side effects would destroy this assumption.

Note that the rule for evaluation of variables does not require a variable to be initialised and, of course, this could cause errors. In order to avoid this problem, all variables can be automatically initialised in the rule for program interpretation. These have been omitted for brevity. An alternative would be to modify the evaluation rule for  $e \in Id$  with an additional premise such as  $e \in \mathbf{dom} \sigma$ .

If a program body is a single statement, this is most usefully a *Compound* (cf. Fig. 1); its interpretation is defined by the interpretation of each of the statements in (left to right) order. The rule for interpretation of *Compound* statements is as follows.

$$\frac{(s, \sigma) \xrightarrow{st} \sigma' \quad (mk-Compound(rest), \sigma') \xrightarrow{st} \sigma''}{(mk-Compound([s] \frown rest), \sigma) \xrightarrow{st} \sigma''}$$

Here the state produced by the interpretation of the first statement,  $s$ , in the list is the state ( $\sigma'$ ) in which to interpret the rest of the statement list, *rest*. As this description is recursive, a base case is required and here this is reached once the list of statements becomes empty. The rule is applied by pattern matching against the input and at this point simply results in an unchanged state.

<sup>9</sup> Technically, the relations  $\xrightarrow{st} / \xrightarrow{ex}$  are the least relations satisfying the rules.

$$\overline{(mk\text{-}Compound([], \sigma) \xrightarrow{st} \sigma)}$$

The SOS rules given so far embody the so-called *big step* operational semantics, as it directly defines the final state. This approach is also referred to as *natural semantics* by Kahn [Kah87] and Nielson and Nielson [NN92]. *Small step* operational semantics has to define the granularity at which interference can occur in concurrency and thus shows the steps between smaller portions of program text and state—the overall interpretation of a program is then the transitive closure of the step relation. Big step tends to feel more intuitive in its handling of multiple statements (and especially constructs like blocks); however, it is worth mentioning the existence of small step concepts because these are used later when concurrency comes into play in Sect. 4.

The core language could be extended to consider some form of external storage such as files with the addition of *Read/Write* statements; this would be accomplished simply by extending  $\Sigma$  to include a collection of (named) files.

## 2.2 Denotational Approach

For simple languages, the difference between the operational and denotational approaches is less marked than when language aspects such as jumps, non-determinacy or the passing of functions as arguments have to be modelled. One important point is that both approaches are built around an explicit notion of state. The technical distinction between operational and denotational approaches is, however, important and the point can be made by contrasting with the earlier *abstract interpreter* phrase: denotational semantics is more like a compiler in that it maps the source language into another language. For the simple language that is defined operationally in Sect. 2.1, the mapping ( $M$ ) would be into functions from states to states ( $\Sigma \rightarrow \Sigma$ ). This state is the same as defined in the previous section.<sup>10</sup> Thus:

$$M : Stmt \rightarrow (\Sigma \rightarrow \Sigma)$$

and the convention of surrounding the (abstract) text parameters by  $\llbracket \cdot \rrbracket$  is followed.

A language is needed to define the functional denotations and Church's Lambda notation is the standard as it provides an easy way to write un-named functions.<sup>11</sup> As a simple example, the assignment statement is mapped to a function which takes a state and returns that state modified with a mapping from the identifier to the evaluation of the right-hand-side expression in the previous state.

$$M\llbracket mk\text{-}Assign(lhs, rhs) \rrbracket = \lambda \sigma \cdot \sigma \uparrow \{lhs \mapsto eval(rhs, \sigma)\}$$

<sup>10</sup> An Oxford denotational semantics would insist that  $\Sigma$  was also a general function type; here the finite, constructed, mappings of VDM are used for  $\Sigma$  because this is not a significant issue in the comparison.

<sup>11</sup> Familiarity with this notation is assumed; a good learning resource is [AGM92].

Much is made in the literature on denotational semantics about the mapping to denotations being *homomorphic* in the sense that the structure of the commands in the object language matches the structure of the denotations. So for compound statements:<sup>12</sup>

$$\begin{aligned} M\llbracket mk\text{-Compound}([\ ])\rrbracket &= \lambda\sigma \cdot \sigma \\ M\llbracket mk\text{-Compound}([s] \curvearrowright rest)\rrbracket &= M\llbracket mk\text{-Compound}(rest)\rrbracket \circ M\llbracket s\rrbracket \end{aligned}$$

Here it can be seen that the sequence concatenation on the left matches the function composition on the right and thus the structure is preserved. The homomorphic property is that the denotation of the compound is built (only) from the denotations of its constituent statements.

Note that the loss of referential transparency requires the state notion. This is now so familiar that it is taken for granted but assignments themselves complicate the denotational ideal of the homomorphic mapping.

It is not difficult to see that there is a clear connection between operational and denotational descriptions (postponing for the moment issues of non-termination):<sup>13</sup>

$$\begin{aligned} interpret : Stmt \times \Sigma &\rightarrow \Sigma \\ M : Stmt &\rightarrow \Sigma \rightarrow \Sigma \end{aligned}$$

$M$  is the Curried form of *interpret*—they are essentially a  $\lambda\sigma$  apart:

$$M\llbracket s\rrbracket = \lambda\sigma \cdot interpret(s, \sigma)$$

But Sect. 2.4 makes clear that the surface difference has a significant impact on reasoning about language descriptions.

The semantics of conditional statements is:

$$\begin{aligned} M\llbracket mk\text{-If}(test, then, else)\rrbracket &= \\ \lambda\sigma \cdot \text{if } M\llbracket test\rrbracket(\sigma) = \mathbf{true} \text{ then } &M\llbracket then\rrbracket(\sigma) \text{ else } M\llbracket else\rrbracket(\sigma) \end{aligned}$$

and again is similar to the operational semantics given in the previous section.

However, the denotational definition of *While*:

$$\begin{aligned} M\llbracket mk\text{-While}(test, body)\rrbracket &= \\ \lambda\sigma \cdot \text{if } M\llbracket test\rrbracket(\sigma) = \mathbf{true} & \\ \text{then } M\llbracket mk\text{-While}(test, body)\rrbracket \circ M\llbracket body\rrbracket & \\ \text{else } \lambda\sigma \cdot \sigma & \end{aligned}$$

includes  $M\llbracket mk\text{-While}(test, body)\rrbracket$  which makes it clear that *fixed points* are required (and this could be made completely explicit by using the fixed point operator  $\mu$ ).<sup>14</sup>

<sup>12</sup> It would be more common to write a denotational description without the constructor (*mk-Compound*) but it has been made clear above that larger languages require an abstract syntax and choosing to keep the same treatment of syntactic objects in the sketched operational and denotational descriptions is useful.

<sup>13</sup> Here, McCarthy's original *interpret*-style description [McC66] is used to make the point more clearly than can be done with the SOS rule.

<sup>14</sup> In early versions of denotational semantics, Christopher Strachey used the  $Y$  combinator to denote the fixed point of a while loop (see for example [Wal67, p. 17]).

### 2.3 Axiomatic Approach

The widest use of *Hoare axioms* [Hoa69] is in the verification or development of programs. It was, however, precisely concerns about ‘leaving things undefined’ in language semantics that led Tony Hoare to propose *Hoare triples*.<sup>15</sup> Perhaps the strongest case for specifying a range of permissible results is in languages that allow concurrent execution and this topic is reviewed in Sect. 4.2. Here, the axiomatic method is explained with the simple sequential language that has been introduced above.

In a deviation from the approach used in the paper so far, concrete syntax will be used in the sections concerned with axiomatic semantics. This is purely by convention: while there is no reason *not* to use abstract syntax, doing so would be unique amongst all other works on axiomatic semantics. The reason for the lack of use of abstract syntax is probably connected to the small scale (and relative syntactic paucity) of the languages to which axiomatic semantics is normally applied.

A so-called *Hoare triple* consists of a pre condition, program text and a post condition. These are now almost universally written as  $\{P\} S \{Q\}$ .<sup>16</sup> In the most widely adopted style, the pre and post conditions are predicates of single states. Note that in contrast with operational and denotational semantics, these states are not explicitly defined. The triple  $\{P\} S \{Q\}$  records a judgement that if  $S$  is executed in a state that satisfies the predicate  $P$ , then (providing  $S$  terminates) the resulting state will satisfy the predicate  $Q$ .

Given this interpretation, inference rules can be provided for each language construct:

$$\begin{array}{c}
 \boxed{\text{Sequence}} \frac{\begin{array}{c} \{P\} S1 \{Q\} \\ \{Q\} S2 \{R\} \end{array}}{\{P\} S1 ; S2 \{R\}} \\
 \boxed{\text{If}} \frac{\begin{array}{c} \{P \vee b\} Th \{Q\} \\ \{P \vee \neg b\} El \{Q\} \end{array}}{\{P\} \text{ if } b \text{ then } Th \text{ else } El \text{ fi } \{Q\}} \\
 \boxed{\text{While}} \frac{\begin{array}{c} \{P \vee b\} S \{P\} \end{array}}{\{P\} \text{ while } b \text{ do } S \text{ od } \{\neg b \vee P\}}
 \end{array}$$

The predicate  $P$  in the rule for **while** is a *loop invariant* and this concept is a key contribution to the way users think about programs even if they are not reasoning completely formally. As noted above, programming constructs can be used to extend what can be expressed in a language. It remains true however that if for example a loop is used to compute factorial, the proof needs axioms about factorial in addition to the inference rule for while statements.

<sup>15</sup> The background to [Hoa69] includes Bob Floyd’s [Flo67] and is traced in [Jon03]; since that publication, earlier drafts have been found of Hoare’s attempts to build on his comment, made at a conference in 1964 [Ste66, pp. 142–143], that “What is required is a method of describing a class of implementation ...”.

<sup>16</sup> In Hoare’s original paper [Hoa69], he actually wrote  $P \{S\} Q$  but placing the braces around the assertions emphasises their role as being non-executable.

The caveat above about termination is important: the *While* rule does not on its own establish that the loop will terminate. This property of correctness assuming termination is often (badly) termed *partial correctness*. Dijkstra [Dij76] proposed the addition of *variant functions* to reason about termination and these were in fact employed without that nomenclature in both [Tur49] and [Flo67]. A more pleasing approach is indicated below when the switch to relational post conditions is discussed.

In practice, users are unlikely to give a post condition in exactly the form  $\neg b \vee P$ . Either the inference rules need to be complemented with weakening rules such as:

$$\boxed{\text{consequence}} \frac{\begin{array}{c} P' \Rightarrow P \\ \{P\} S \{Q\} \\ Q \Rightarrow Q' \end{array}}{\{P'\} S \{Q'\}}$$

or, perhaps more usefully, the other rules should be changed to reflect the potential for weakening—for example:

$$\boxed{\text{While}' } \frac{\begin{array}{c} \{P'\} S \{P\} \\ P \vee b \Rightarrow P' \\ P \vee \neg b \Rightarrow Q \end{array}}{\{P\} \mathbf{while} \ b \ \mathbf{do} \ S \ \mathbf{od} \ \{Q\}}$$

Having considered the sort of statement that controls the order in which basic statements are executed, the axiomatic description of assignment statements must be addressed. The now standard<sup>17</sup> *backwards rule* can be written

$$\boxed{\text{assign}} \frac{}{\{P_e^x\} x := e \{P\}}$$

where  $P_e^x$  means substitution of  $e$  for  $x$  (with appropriate renaming to avoid unwanted capture).

The deceptively simple—and therefore appealing—rule is not without its problems. For example Krzysztof Apt in [Apt81] discusses the careful adjustments required if the left-hand-side of the assignment is a reference to an element of an array. Without wishing to undervalue what might be thought of as a lucky notational success, it must be observed that the aforementioned lack of referential transparency with variables in programs should prompt care when copying their names into predicates.

Another reservation about the assignment rule arises when languages allow multiple identifiers to refer to the same location (see Sect. 3.3); sticking to the

<sup>17</sup> Floyd in [Flo67] used a forward assignment axiom that needs an existential quantifier in its post condition; having discussed the developments with several people (including King whose Effigy system [Kin69] used the backward rule) it would appear to be the case that Bob Floyd spotted the simpler rule after his paper was published and that David Cooper took the information from Carnegie Tech (where he had been for over a year) to Tony Hoare in Belfast when Cooper gave a seminar there.

assignment rule above would appear to imply that *call-by-reference* is modelled by some form of copy rule.<sup>18</sup>

In [Hoa69], Tony Hoare indicates that the axiomatic approach obviates the need for an explicit model of state. This connects with the well-known *frame problem* in the sense that it would be convenient if the only thing affected by an assignment to  $x$  is the value of the variable with that name. This is, of course, not the case in the presence of call-by-reference parameter passing.

It was realised early on<sup>19</sup> that writing relatively large collections of axioms could lead to inconsistencies. The standard way out of this danger is to provide a model for which axioms can be shown to hold. Under Tony Hoare's supervision, this is exactly what Peter Lauer undertook in his thesis [Lau71]; a later—but better-known—reference is [Don76]. Essentially, it is necessary to show that if  $\{P\} S \{Q\}$  can be deduced from the axioms, then this agrees with the operational semantics as follows:

$$P(\sigma) \vee ((S, \sigma) \xrightarrow{st} \sigma') \Rightarrow Q(\sigma')$$

If termination is considered, it is also necessary to show:

$$P(\sigma) \Rightarrow \exists \sigma' (S, \sigma) \xrightarrow{st} \sigma'$$

The *sequence* axiom above shows clearly why it is attractive to use post conditions that are predicates of a single state. It should, however, be obvious that this is not really a good idea! What a program is intended to realise is a final state that relates in some meaningful way to its initial state. VDM has used relational post conditions since before [Jon80]—Aczel showed in an unpublished note [Acz82] how to present rules for such relational specifications in a convenient way—and these rules of inference are used in subsequent VDM publications. A particular advantage of explicitly using relations is that Dijkstra's *variant functions* are avoided simply by saying that the specification of the body of a loop should be a well-founded relation.

Hoare's 1969 paper is one of the most influential references in theoretical computer science. It can be seen as the root of developments including Edsger Dijkstra's *weakest pre conditions* [DS90] and work on *refinement calculus* [Mor94, BvW98]. Furthermore, this whole line of thought led, after [Hoa71b], to the use of Floyd/Hoare axioms in the development process (rather than *post facto* proof). Further discussion of these developments is available in [Jon03].

## 2.4 Reasoning

There are two distinct needs to reason based on a (formal) semantics. On the one hand, a programmer might want to prove that a program satisfies its specification; on the other, the designer of a compiler might want to justify the design

<sup>18</sup> Various other extensions by Hoare include [CH72, Hoa72a, Hoa71a]; useful summaries are [Apt81, Apt84].

<sup>19</sup> Specifically at the April 1969 IFIP WG 2.2 meeting in Vienna at which Hoare first presented his axiomatic method [Wal69]. See [JA16] for more comments on this meeting.

of a compiling algorithm. (In both cases, the more important issue is how to use the semantics as the basis for a stepwise development but that does not affect the distinction.) Here, both tasks are first explained in terms of operational semantics.

In proving the correctness of a program, its specification should take the form of a pre condition and a post condition. The first of these describes any assumptions on the state before execution of the program; the second defines the acceptability of the state produced after the program as a relation to the initial state. The post condition is a predicate of *two* states (before and after) because all but the most trivial specifications relate values in the post-state to those in the pre-state (as with defining the result of a function with respect to its arguments):

$$\begin{aligned} pre &: \Sigma \rightarrow \mathbb{B} \\ post &: \Sigma \times \Sigma \rightarrow \mathbb{B} \end{aligned}$$

This specification is related to the implementation by formulating the related *Proof Obligation* for the program  $S$ :

$$\forall \sigma \in \Sigma \cdot pre(\sigma) \Rightarrow post(\sigma, interpret(S, \sigma))$$

Discharging this proof obligation indicates that the program  $S$  satisfies the specification given.

In the task of proving correctness of translation, the proposed algorithm might have the signature:

$$translate : Stmt \rightarrow MachineCode$$

and the machine code might be given semantics by:<sup>20</sup>

$$mc\_interpret : MachineCode \times \Sigma \rightarrow \Sigma$$

This allows us to formulate the proof obligation as follows:

$$\forall S \in Stmt, \sigma \in \Sigma \cdot mc\_interpret(translate(S), \sigma) = interpret(S, \sigma)$$

Although it is possible to reason about the earlier program correctness task using either an operational<sup>21</sup> or a denotational language description, that is exactly the task for which axiomatic semantics was envisioned.<sup>22</sup>

In contrast, the task of reasoning about the correctness of a language translator appears to be best handled with one of the *model oriented* (i.e. operational or denotational) description methods. The choice between operational and denotational semantics as a basis for such proofs depends on a number of factors. The higher level of abstraction in noting that denotations are functions (for now, from states to states) certainly makes it easy to establish some properties

<sup>20</sup> This has been deliberately simplified by ignoring the fact that the abstract states ( $\Sigma$ ) of the language description need to be reified to representations on the object-time storage organisation.

<sup>21</sup> This approach is explored in John Hughes' thesis [Hug11] and [HJ08].

<sup>22</sup> As observed in Sect. 2.3, such proofs also require axioms of any new operators.

of a language (e.g. the equivalence of a while loop to its unwrapping with a conditional around the original loop).

For translation algorithms that closely follow the phrase structure of the source language, denotational semantics is probably most appropriate because it is easy to reason about the functional semantic objects. Robert Milne and Christopher Strachey tackle implementation correctness in both the “Adams Essay” [MS74] and the two-volume book [MS76] published after Strachey’s death; members of the IBM Lab Vienna addressed compiler correctness using denotational semantics as well. Unfortunately, as the latter were concerned with the large (and Baroque) language PL/I, most of the publications are only available as lengthy technical reports (e.g. [BBH+74, Wei75, Izb75, BIJW75, Jon76]).

Unfortunately, many compiling techniques are not obviously algebraic in form: optimisations such as register allocation or *strength reduction*<sup>23</sup> cut right across the phrase structure of the language and cause problems for descriptions reliant on homomorphic denotations. In such cases, it might well be easier to base the argument on an operational description—publications on using operational descriptions to reason about compiling include [MP67, Pai67, Luc68, Jon69, JL71].

One point of comparison that is worth clarifying is that operational semantics can be made to appear as compositional as denotational semantics. It is true that early attempts to provide operational semantics of large languages (e.g. the VDL descriptions of PL/I [WAB+68]) often fell into the trap of putting things in the state that were unchanged by simple statements—McCarthy referred to this as the *grand state* mistake. Furthermore, seeking a homomorphic mapping (to functional denotations<sup>24</sup>) encourages someone writing a denotational semantics to consider exactly what must be in the state. But a small state SOS description can closely follow the phrase structure of the language being described. The main penalty for using, for example, an SOS description is that proofs have to use induction over the steps of computation rather than, say, Scott induction [Win93, p. 166].

One significant point in the comparison of denotational and operational descriptions concerns termination. The program

**while**  $x \neq 0$  **do**  $x := x - 1$  **od**

will, for a negative initial value of  $x$ , simply iterate indefinitely. Reading a big step (or *natural*) operational semantics as inference rules means that the hypotheses cannot be discharged for such values. In contrast, the least fixed point of the denotation of this program is exactly the partial function that takes states with positive values for  $x$  to states where  $x = 0$ .

The greatest payoff for the level of abstraction in denotational semantics is in proving deeper properties of a defined language.

<sup>23</sup> Within a loop, a relatively expensive operation such as multiplication can be replaced by addition if one of the operands is the control variable of the loop.

<sup>24</sup> Finding neat functional denotations is not always possible. The topic of abnormal exits such as **goto** statements is postponed to Sect. 6 but forces considerable contortions of the space of denotations.



## 2.5 Section Summary

The main challenge presented by simple imperative languages is the need to store and update values associated with variables when assignments are made. The response given by both operational and denotational semantics is to model the storage of the computer with an abstract state. There is no fundamental difference between the states used in denotational and operational semantics. Axiomatic semantics avoids an overt state by using value replacement, but the collection of meta-variables used in assertions does essentially imply a state.

## 3 ALGOL-Like Blocks, Functions, Procedures

For the simple language presented above, the differences between the semantic description styles seem minor. But that language lacks many features that make real languages convenient for programmers. The challenge of describing language features like named procedures and environmentally-separated blocks adds significant complexity to the task of language description and begins to show interesting differences in the response by each semantic school.

The need to model the local entities of different blocks and sharing of locations presents particular challenges, especially in the presence of more complicated data structures such as arrays. Procedures add additional problems when different parameter mechanisms are considered and so-called *higher-order* procedures (whose parameters or results are procedures themselves) are particularly problematic in some approaches. This section discusses these challenges and the solutions in the different approaches.

It is interesting to observe how similar the treatment is in denotational and operational approaches—and to note the key difference on procedure denotations. Axiomatic semantics ends up taking a different tack by avoiding environments and instead using name substitution.

### 3.1 Local Naming

In first-order predicate calculus:

$$\forall x \in X \cdot (\dots \forall x \in Y \cdot (\dots) \dots \exists x \in Z \cdot (\dots))$$

the three bindings of  $x$  are distinct: they occupy separate *name spaces*. The need for separate name spaces in programming languages is even stronger because program texts are likely to be long.

Most programming languages offer ways of localising a name space so that the same identifier can denote a different variable in nested *blocks*.

$$Stmt = \dots | Block$$

$$\begin{aligned} Block &:: types : Id \xrightarrow{m} ScalarType \\ &\quad body : Stmt \\ &\quad \dots \end{aligned}$$

In the simple storage model of Sect. 2, identifiers are mapped to denotations (so far only values) and there is no need so far to change the underlying state notion. The only delicate point is that – at block exit – the semantics must recover the denotations of those identifiers that referred to a different variable in the inner block.

Context conditions must also be reconsidered now that the same identifier may denote different values and types throughout computation. This can be achieved by requiring that usage of names in a well-formed block matches the closest embracing declaration. A well-formed program now need only require that every constituent block is well-formed.

### 3.2 Functions, Procedures and (Simple) Parameters

The pragmatics of functions and procedures is that they can be used to factor out portions of program text that can be called from many places.<sup>25</sup> From a user point of view, procedure calls are statements that get executed in the order dictated by their position in a list of statements whereas functions occur in expressions.<sup>26</sup> Functions and procedures require similar modelling techniques in terms of the semantic objects required and are therefore treated together in the remainder of this section.

$$\begin{array}{l} \textit{Block} :: \dots \\ \qquad \textit{body} : \textit{Stmt} \end{array}$$

$$\textit{Stmt} = \dots | \textit{ProcCall}$$

Context conditions of procedures are similar to those for blocks, but additionally require that the evaluated types of parameters in a procedure call match those declared in the procedure definition. This means that the *TypeMap* object must also store information on procedure definitions.

Functions and procedures have fixed denotations so they do not belong in the store which contains values that can be changed (by assignment) within statements. This can be handled by introducing an *environment* to contain the denotations:

<sup>25</sup> Although compiling techniques are not the main topic of this paper, it is worth observing that implementing general recursion and parameter passing required the invention of ingenious techniques—see [RR62]; there is a very detailed reconstruction of the development of the idea of the *Display* mechanism in [vdH17].

<sup>26</sup> It is worth noting that functions which can cause side effects considerably complicate expression evaluation. At a minimum, they remove the possibility of saying that  $\textit{eval} : \textit{Expr} \times \Sigma \rightarrow \textit{Value}$  because of the potential state change inherent if functions with side effects are allowed within *Expr*. Something that causes language descriptions more trouble is that, unless the order of evaluation of expressions is strictly defined (which is rare because languages tend to leave compilers the freedom to optimise register use), evaluating expressions containing functions with side effects results in non-determinism. This general topic is resumed in Sect. 4.

$$Env = Id \xrightarrow{m} Den$$

$$Den = FunDen | ProcDen | \dots$$

The basic model is not difficult; that having been said, the features that have been devised in various languages to make procedures more useful are myriad and necessitate extension of the role of the environment. The passing of parameters of simple values (e.g.  $\mathbb{N}, \mathbb{B}$ ) is straightforward: these are simply given new identifiers within the local environment of the function or procedure. However, more complex parameter passing mechanisms require more consideration.

### 3.3 Sharing

Thus far, it has been assumed that identifiers denote simple distinct values such as numbers or Booleans. However, for reasons of efficiency, it is sometimes useful to have more than one identifier referring to the same entity. Because of potential name clashes, making precise the semantics of such sharing is non-trivial. Classically, logicians (e.g. in describing the Lambda calculus) have used a *copy rule* with “suitable changes of names to avoid clashes” to describe such concepts. For programming languages, the text of the procedure can be modified to copy in the names, references or values of arguments, with appropriate renaming to avoid name clashing. The ALGOL report [BBG+60] uses an informal description of this approach to attempt to fix the semantics; it can also be formalised, as in the operational description of ALGOL 60 [ACJ72].

Many programming tasks require composite entities such as arrays which gives rise to the notion of *left hand values* for elements of arrays. These considerations are the main reasons for allowing different ways of passing arguments to functions or procedures. Surprisingly many alternative parameter passing mechanisms have been devised and each has its use:

- Call *by value* is the most obvious and is appropriate for simple types—the argument (which might be an expression) is evaluated and this value is copied into the body of the function or procedure. Typically this is achieved by creating a new memory allocation for the value and therefore modifications to this variable are not seen in the calling scope.
- Copying of data can be reduced by using *by location* (or *by reference*) parameter passing, in which a pointer to the storage location of the argument is passed instead of its value. This enables the function to modify the value of the argument variable in a way that will affect the calling context.
- The full *by name* parameter mechanism of ALGOL 60 is even more challenging semantically: the denotations of arguments are evaluated anew each time the respective parameter name is encountered within the body of the function, thereby potentially triggering multiple instances of side-effects. (This specialises to *by location* mode when the argument (or *actual parameter* in ALGOL speak) is a simple identifier.)

- Call by *value/result* offers a useful compromise; by copying the value of each argument into a new location and then returning the (potentially modified) values to their original locations; it facilitates the return of multiple values from procedures/functions but avoids the problem of the same location being referred to by different identifiers.<sup>27</sup>

In model-oriented methods, all of the above can be modelled with:<sup>28</sup>

$$Env = Id \xrightarrow{m} Den$$

$$Den = \dots | Loc$$

$$Loc = ScalarLoc | ArrayLoc$$

$$ArrayLoc = \mathbb{N}^* \xrightarrow{m} ScalarLoc$$

$$\Sigma = ScalarLoc \xrightarrow{m} ScalarValue$$

In SOS it is clear that the environment is not changed by simple statements such as assignments as *env* is not in the range of the  $\xrightarrow{st}$  relation.

$$\frac{(rhs, env, \sigma) \xrightarrow{ex} v}{(mk - Assign(lhs, rhs), env, \sigma) \xrightarrow{st} \sigma \uparrow \{env(lhs) \mapsto v\}}$$

The task of creating and passing locations is handled in the semantics of blocks and calling.

Similarly, in denotational semantics, the fact that environments are not changed by simple statements is apparent from the *Curried*:

$$M : Stmt \rightarrow Env \rightarrow \Sigma \rightarrow \Sigma$$

It is interesting the extent to which the description of semantic objects and a few type definitions (i.e. no actual rules or formulae) can suggest (to an experienced reader) the main points about a language. The rest of this paper is written at this level of abstraction.

The passing of parameters in environment-based semantics is not difficult—the semantic function, relation or mapping is extended to include an environment as a parameter and this environment is modified at evaluation time. The parameter passing mechanism chosen affects the level at which the environment or its sub-contents are modified.

It is, however, important to clarify how the context of a procedure or function is captured in model-oriented approaches. In an operational approach, one part of *ProcDen/FunDen* is its text. But this is not enough: if functions/procedures can be declared in any block and called from any deeper block, then there must be a way of fixing the *environment* in which they are to be executed, so that

<sup>27</sup> Unless the same argument is passed to different parameters—but this is an easy static check.

<sup>28</sup> Records are similar to arrays but have fields that are not necessarily of the same type; modelling records and combinations of arrays/records is straightforward.

there is a proper evaluation of any parameter identifier that is passed in, and no clashes with local names used within the text of the procedure. To address this, an environment is usually part of the interpreting function or relation for procedures and functions. This approach is essentially identical to the *static chain* method for address resolution, in which each scope contains some meta-information linking it to its direct lexical parent.

In denotational approaches, *FunDen/ProcDen* are functions in the standard mathematical sense, with the appropriate environment bound in forming a *closure*.<sup>29</sup> Environments are therefore also parameters to the meaning function, as seen above.

### 3.4 Handling Parameters and Sharing in the Axiomatic Approach

Using *by location* parameter passing means that multiple identifiers refer to the same *location* and, at a minimum, this undermines the axiom of assignment in Hoare triples. So the axiomatic approach, tending to ignore the concepts of both state and environment, uses quite a different strategy to model-oriented techniques: a form of repeated name substitution is used, essentially a modification of the copy rule described above.

The basic case for the invocation of a procedure is one where there are no parameters and no side effects; calling a procedure is essentially adding the body of the procedure to the main program body. The following simple rule (adapted from [Pag81]) applies:

$$\boxed{\text{Invocation}} \frac{\begin{array}{l} \{P\} S \{Q\} \\ N.\text{body} = S \end{array}}{\{P\} \text{ call } N \{Q\}}$$

Adding parameters requires that variables in  $P$  and  $Q$  referring to the parameters of  $N$  be replaced by the arguments (or argument expressions, or evaluated argument expressions, depending on calling mechanism). Such substitution must be conflict avoiding, but this is just generally assumed to be taking place rather than explicitly mechanised in axiomatic descriptions.

$$\boxed{\text{Invocation}'} \frac{\begin{array}{l} \{P\} S \{Q\} \\ N.\text{body} = S \\ N.\text{params} = [N_1, \dots, N_n] \end{array}}{\{P_{E_1, \dots, E_n}^{N_1, \dots, N_n}\} \text{ call } N(E_1, \dots, E_n) \{Q_{E_1, \dots, E_n}^{N_1, \dots, N_n}\}}$$

Procedures with side effects can also be handled, and a way is provided in the (incomplete) axiomatic ‘definition’ of Pascal. This approach expands the notion of parameters to include all variables used globally within  $N$  and considers these to be ‘implicit’ parameters. They are then handled in the same way as ‘explicit’

<sup>29</sup> As is the case with axiomatic semantics in Sect. 2.3, strictly, the function itself is not produced: the semantics maps to a Lambda expression that could be proved equivalent to the mathematical function using properties about the function.

parameters: functions are assumed to exist which map the initial values of both explicit and implicit parameters onto their final values and these are used in the assertion substitutions as in the rule *Invocation'* above.

Arrays (even without sharing) need careful handling in axiomatic semantics, as also discussed by [Apt81]. Allowing expressions as the subscripts in subscripted variables can lead to problems, particularly when these expressions reference the same array. One way to address this is to replace the whole array with a new one modified at the index to which assignment has been made, but this is not a particularly elegant solution.

### 3.5 Higher-Order Functions and Procedures

The pragmatics of allowing parameters to be procedures and functions is to facilitate higher-order programs. Not only is this concept beloved by functional language users, it is also a prime tool for abstraction in programming. For example, the simple *map list* idea

$$\text{map-list} : (A \rightarrow B) \times A^* \rightarrow B^*$$

provides a generic function that yields a sequence in which every element is the result of applying the function in the first argument to the corresponding element of the second argument; this is a small example of how high levels of re-use and abstraction can be achieved. There are, of course, far more exotic cases that introduce new ways of achieving recursion: see, for example, Knuth's "man and boy" example [Knu64] that was written as a challenge for ALGOL 60 compilers.

This topic is placed in a separate sub-section because it causes one of the most telling differences between operational and denotational approaches. The clue to the source of the problem is that, once functions can take functions as arguments, the possibility arises that a function can be applied to itself. (This also introduces a minor issue around types that is reviewed at the end of this sub-section.)

The fact that, in operational semantics approaches, the denotation of a procedure is a pair (containing the text of the procedure and its statically containing environment) means that no new concepts are needed to model the passing of procedures or functions.

In denotational approaches, however, the denotation of procedures are actual functions (as indicated in Sect. 3.3). During the development of denotational semantics, this gave rise to a serious mathematical problem: since the cardinality of the function space  $X \rightarrow X$  must be greater than that of  $X$ , there is a paradox with functions that can take themselves as arguments. There was thus a point in time where Strachey's idea of *denotational* (or at that time *mathematical*) semantics claimed that semantics could be given by mapping programs to mathematical functions (expressed in the Lambda calculus), but the approach was built on sand in the sense that no one could offer a model of the untyped Lambda calculus.

This problem was resolved with Dana Scott's 1969 invention of domains with suitably restricted functions. This was a major intellectual achievement and has been widely described; perhaps the most accessible text remains [Sto77] but Scott's own [Sco80] provides a clear description of the context of his models of the untyped Lambda calculus.

The challenge of modelling self-applying functions gives rise to the largest divergence so far between operational and denotational approaches. It is interesting to look more carefully at what is going on here. The *homomorphic* rule says that the denotation of a construct should be built up from the denotations of its constituent parts. But the name of a procedure can only be given a denotation by storing it in an environment.

There is, in fact, another issue to be resolved for functions that can take themselves as arguments; that issue concerns defining their type. Consider first a binary tree structure built up with records:

$$\begin{array}{ll} \text{BinTree} :: & \text{left} : [\text{BinTree}] \\ & \text{value} : \mathbb{N} \\ & \text{right} : [\text{BinTree}] \end{array}$$

The name of the type *BinTree* is used to express the recursive embeddings and the marking of the fields as optional ensures that instances can be finite.

In order to declare a function type that can take itself as argument, there must be a way of naming a function type. In fact, ALGOL 60 ducked this problem: the language is almost strongly typed except for function and array types. Both PL/I and Pascal offer such separate naming of function (entry) types. It is worth noting that separating function types is necessary for mutually recursive procedures because they cannot be given in an order such that each definition precedes use.

### 3.6 Section Summary

Blocks and procedures bring new challenges to semantic descriptions, particularly with the concerns of name sharing and local entities. Denotational and operational semantics solve this problem by separating out an environment from the state, but very cautious name substitution is needed in axiomatic semantics, particularly when advanced parameter mechanisms are used. Procedures become another kind of denotable value in model based semantics, but this requires careful foundation for denotational semantics when higher-order functions are allowed.

## 4 Modelling Non-deterministic Languages

There are two essentially different reasons that non-determinism figures in programming languages:<sup>30</sup>

<sup>30</sup> A separate need to have a formal treatment of non-deterministic specifications arises when considering *program development*—see Sect. 4.2.

- the originator of a language might wish to allow freedom to the designers of implementations to make optimisations such as common sub-expression elimination;
- a language might encompass features that result in non-deterministic execution—the most telling example is concurrency where differing progress of threads can yield a range of results for executing a program.

It is clear that the specification (or description) of a language must fix the full—and exact—range of acceptable outcomes. This matters both to programmers writing programs in the language and language implementers. The challenge is leaving some aspects of the language incompletely defined, but properly constrained. This problem is further complicated by questions of *granularity* of interleaving: a semantic description must be capable of describing granularity at least as fine as that handled by the language. The difficulty of these points is a significant challenge for the semantic description: having a sufficiently rich notation to allow communication of these aspects while remaining readable. These challenges existed as soon as languages such as PL/I were addressed; the various responses are interestingly different in appearance but do have a common core.

#### 4.1 Operational Response

The pragmatics of concurrent programming languages should be obvious: both low-level systems programming and high-level applications need to express algorithms that accommodate differing run-time progress. In model-oriented semantic approaches, there appears to be no alternative to recording the text of the threads that remain to be executed and adjoining it to the shared state ( $\Sigma$ ) that is being updated. Such pairings of states and remaining thread texts are referred to as *configurations*.

In order to capture the possible mergings of the threads, an operational semantics must show the non-deterministic choice between the threads. Precisely how this is done fixes the granularity of merging.<sup>31</sup> A first thought might be to record a function that maps a configuration to the set of its possible successor configurations but this becomes notationally messy. It is, of course, equivalent to think of this as a relation between configurations and it transpires that this is notationally much cleaner. There are many ways to define such a relation. The approach utilised in the early operational semantics VDL documents [Lab66, LW69], offered a way of describing such non-determinacy by using *control trees* that contain a structured version of the program text that still had to be executed—but these control trees were made part of the (grand) state.<sup>32</sup> Plotkin's SOS [Plo81] provides much clearer descriptions because the

<sup>31</sup> Many attempts to provide ways of reasoning about concurrent programs (see Sect. 4.2) make the assumption that assignment statements are atomic; for brevity, this simplification is followed here; but it must be realised that this level of granularity is unrealistic for real implementations of languages due to the possibility of values of variables being changed by parallel threads even during expression evaluation.

<sup>32</sup> For a fuller discussion see [JA16, Sect. 3].



non-determinacy is factored out of the rules themselves; it moves to the selection of a semantic rule (the remaining text and state are kept separate).

With the following definition of *Parallel* consisting of two threads

$$Parallel = (Thread \times Thread)$$

$$Thread = Assign^*$$

a large-step approach is inappropriate: an interpreting rule like  $\xrightarrow{st}$  from Sect. 2.1 would interpret an entire sequence of assignments as one. This limits the language to executing the *Parallel* as though each *Thread* were atomic. What is needed is a set of rules which each peel off and execute one of the remaining statements in any non-empty thread. For this we use the relation for parallel interpretation,  $\xrightarrow{par}$ . A *small step* semantics interprets the next assignment in either the left or right thread:

$$\xrightarrow{par}: \mathcal{P}((Parallel \times \Sigma) \times (Parallel \times \Sigma))$$

$$\frac{(s, \sigma) \xrightarrow{st} \sigma'}{([s] \curvearrowright restl, r), \sigma \xrightarrow{par} ((restl, r), \sigma')}$$

$$\frac{(s, \sigma) \xrightarrow{st} \sigma'}{(l, [s] \curvearrowright restr), \sigma \xrightarrow{par} ((l, restr), \sigma')}$$

Using this approach, assignments may be interleaved in any order, as the choice of which thread to interpret next is lifted to the choice of rule instantiation.

Extensions for other language features can be made in a similar style to this; for example, a small step model of a *while* loop unwraps the loop with a conditional surrounding it.

Note that so far the assumption is that assignment statements represent the level of atomicity in the language. Allowing interference to take place at the expression evaluation level is possible and makes two things clear:

- The way that SOS factors out the non-deterministic choice of rules that match the current configuration is extremely helpful in preventing the issue of concurrency from polluting a whole definition.<sup>33</sup>
- A further observation is that, in SOS descriptions, the non-determinacy with expressions looks different from that with statements: with expressions, the non-determinacy is resolved when a variable is accessed (or a function returns a value) and the effect is to place a value in the evaluation tree; with statements, the effect is reflected in a state change and the executed statement is discarded from the resulting configuration.

<sup>33</sup> But there is a sense in which the configurations are just a way of presenting the *control trees* that were much criticised in VDL operational descriptions (The danger with these control trees in a grand state semantics was that it was hard to determine where they could or could not be updated.).

Moving to a level of granularity larger than assignments, a programmer may wish to make *multiple* statements executable only as an atomic block.

$$Stmt = \dots | Atomic$$

$$Atomic :: Assign^*$$

$$\frac{(sl, \sigma) \xrightarrow{st} \sigma'}{([mk-Atomic(sl)] \curvearrowright rest, \sigma) \xrightarrow{st} (rest, \sigma')}$$

Atomicity is, of course, a key issue in the database world and it is interesting to note the similarities to—and differences from—the programming language universe. It would not be difficult to add data types to a programming language that provide ways to declare and manipulate relations similar to those in the standard relational model (see [Dat82]). As discussed at a Schloss Dagstuhl event on atomicity [JLRW05, Sect. 2.4.2], this then highlights the point that database systems strive to prevent data races, where possible, by system-induced locking (and, where pre-planning fails, to detect races and handle the recovery) whereas programmers using typical programming languages are held responsible to plan and control locking.

## 4.2 Axiomatic Response

As indicated in Sect. 2.3, the axiomatic approach copes with general non-determinism naturally. This observation that it is important to leave aspects of a language undefined was made by Tony Hoare in [Ste66, pp. 142–143] and—via multiple drafts—led him to his famous *axiomatic basis* paper [Hoa69].<sup>34</sup> Moreover, it became clear in using methods such as VDM that specifications that allow a range of implementations are a powerful way of structuring design decisions (see for example [Jon90, Abr10]).

Unfortunately the specific case of non-determinacy being caused by concurrent execution presents severe challenges for the axiomatic approach. The source of the difficulty is precisely the *interference* that has to be modelled explicitly in the operational descriptions of the previous sub-section. Before facing the fact that post conditions alone are insufficient to specify components that suffer interference, it is interesting to trace an early attempt to finesse that difficulty and its more recent manifestation in (Concurrent) Separation Logic.

Hoare singled out the case of disjoint concurrency in [Hoa72b] and made the observation that the post conditions of two parallel threads could be conjoined providing there were no shared variables. Hoare’s 1972 paper covered normal (stack) variables in which case the disjointness is simply a check of the *alphabets* of the threads. John Reynolds introduced *Separation Logic* [Rey78, Rey89] to support reasoning about *heap variables* (i.e. data structures that contain pointers and whose topology can be changed by updating said pointers). Reasoning

<sup>34</sup> Of course, the soundness notion at the end of Sect. 2.3 needs to be enriched but this is straightforward.

about parallel threads that share a heap can be very delicate. An interesting collaborative attack (see [BO16]) led to *Concurrent Separation Logic* [O’H07] which has spawned many variants—see [Par10]. The essential idea is akin to Hoare’s observation: what one wants to do is to conjoin the post conditions of parallel threads but this is only valid if the interference is avoided. What separation logics facilitate is concise statements of the disjoint ownership of heap addresses.<sup>35</sup> More recently, [JY15] notes that certain cases of heap separation can be viewed as reifications of abstract descriptions of separate entities.

In [O’H07], it is suggested that separation logic should be used to reason about race-free programs and *Rely/Guarantee* (R/G) conditions should be used for *racey* programs.<sup>36</sup> The initial publications on R/G go back to [Jon81]—more recently the same underlying concept has been expressed in a refinement calculus [Mor94, BvW98] style in [HJC14, JHC15]. This, in particular, makes algebraic properties such as the distribution of rely and guarantee conditions over sequential and parallel program operators much clearer.

The basic R/G idea is that acceptable interference should be documented with rely conditions in the same way that sequential Floyd/Hoare logic records acceptable starting states with pre conditions. Also, just as post conditions express obligations on the running code, guarantee conditions record the upper limit of interference that a component can inflict on its environment. Specifications of components using R/G conditions can then be used as a basis for design justification. In a step where the sub-components are also specified using R/G conditions, clear proof obligations exist to justify development steps for parallel operators. Unsurprisingly, these proof obligations are more complicated than those for sequential Floyd/Hoare logic but the essential property of *compositionality* is preserved.

Just as at the end of Sect. 2.3 the soundness of these inference rules needs to be proved. It is possible to extend the operational semantics to carry an interference relation and then to interpose it at points appropriate to the granularity of the language; this approach is used in [CJ06, Col08]. Alternatively, *Aczel traces* (see [Acz83] or the more accessible [dR01]) can provide a space of denotations and [CHM16] does this in a way that conducts proofs at a significantly higher level of abstraction.

Another method for modelling concurrency is that of *process calculi* or *process algebras*, which include ACP [BK84], CSP [Hoa85], CCS [Mil89] and  $\pi$ -calculus [SW01]. CSP is particularly relevant due to its influence on the programming

<sup>35</sup> This led Jones to make a suggestion at the MFPS meeting in 2005 where O’Hearn presented concurrent separation logic that it might better be thought of as *ownership logic*.

<sup>36</sup> Although this seemingly simple dichotomy ignores the way in which non-interference at an abstract level can be used to establish race freedom in a representation—a nice example is Simpson’s *Four-Slot* implementation of Asynchronous Communication Mechanisms in [JP11]; this paper also introduced the idea of a notation for *possible values* which is, in turn, explored in [JH16].

language occam, used extensively by Inmos [INM88]. Although work on these approaches grew out of considerations of language semantics, they are no longer strictly within the scope of this paper.

### 4.3 Denotational Response

The key to the utility of a denotational semantic description is the choice of a space of denotations which admit tractable reasoning. Denotations for the language of threads above could be either relations over states or functions from states to sets of states. In either case, there is a need to mark (potential) non-termination. It is important to note that the problem of interference remains: just as an operational semantics must indicate the granularity of thread switching by the way in which configurations are changed and rematched, the relations must be composed appropriately.

Thus far, there is a lot of similarity between denotational and operational presentations of the semantics for non-determinacy resulting from concurrent threads. The combination of non-determinism with higher order functions (cf. Sect. 3.5) however poses extra difficulties for the denotational approach. Here *Power Domains* [Plo76, Smy76] are required to preserve the mathematical properties that overcome the cardinality paradoxes related to higher-order language constructs. Again operational semantics is inherently simpler because procedures and functions are modelled simply by their texts.

### 4.4 Section Summary

The challenges of parallelism bring some variance in the response from the various semantics. In operational semantics, the non-determinism is lifted to the rule level and the real power of SOS to merely constrain acceptable solutions (rather than generate a unique solution) is displayed. In some ways this is similar to certain axiomatic responses, where interference and interaction is constrained by logical propositions. Denotational semantics runs into foundational technicality since the traditional function can no longer be used as a base for denotations. Instead, contortions of the semantic domains such as power domains are required.

## 5 Applying the Ideas to a Concurrent Object-Based Language

This section outlines the semantics of a concurrent object-orient language known as COOL,<sup>37</sup> designed to be small enough to model in a small document but realistic in its handling of the issues identified above.

---

<sup>37</sup> COOL was inspired by – and is similar to – POOL [AR92]. COOL is used in teaching a course on language semantics at Newcastle University.

SIMULA 67 [DMN68] was designed as a language in which simulation programs could be constructed; this provides a wonderful intuition for *Object-Oriented* (OO) programming languages: objects are blocks that can be instantiated as required,<sup>38</sup> block descriptions are the class definitions, local variables are the instance variables and procedures are methods. The scope of method names is of course external to the class to enable objects to call methods defined for other objects.<sup>39</sup>

Key issues in the design of a concurrent language are how to generate and synchronise concurrent threads. Although it gives an unconventional OO language, the aims of this section can be achieved by limiting (instances of) objects to running one method at a time and generating concurrency by arranging that many objects can be active. This ensures that instance variables are free from *data races* and, crucially, that the level of interference is in the hands of the programmer because only by sharing references (to objects) is interference possible.

The move from the unconstrained concurrency of threads in Sect. 4 to a simple OO-language can be summarised as follows:

- The language in Sect. 4 has dangerous data races because of the single shared state.
- In COOL each object (instance) has a local state and can run as a thread.
- Such extreme separation needs to be tempered by providing some communication between the threads. This is easy to achieve by allowing methods to be called in objects. Parameter passing is by value; object references can be passed thus opening up both (controlled) sharing and passing of the ability to invoke methods.
- Any object can create an object (that is an instance of a class) and receives the unique *Reference* of the new object. The relevant statement might be called *New*.
- The only way in which objects can begin execution is by having their methods called by other objects (the exception is for the initial object which begins execution at program start). Objects retain references to their client objects and should eventually cease execution and return values.
- Thus far, there is no obvious source of the claimed concurrency but there are many ways to create parallel threads:
  - A class could have a designated initial method that begins to execute in any newly created object of that class: instantiating multiple objects results in concurrent execution. Similarly, a program could have a set of designated objects which all begin execution when the program starts (this latter approach is presented in the language description below).

<sup>38</sup> When Ole-Johan Dahl made this comment to Jones, the whole OO area became clearer.

<sup>39</sup> The desire to add some notion of object orientation to languages such as C did not necessarily result in languages with clear semantics. SmallTalk [GR83], however, is a principled OO language and Bertrand Meyer's Eiffel language [Mey88] adopts the pre/post specification idea to provide *contracts*.

- ABCL [Yon90] included a *FutureCall* statement that essentially forks the called method—the join occurs when the client object executes a *Wait* statement.
  - An alternative explored in [San99] is to have a *Release* statement that prematurely releases the client object before the server method is complete. Using this strategy, the client can resume execution while the server continues to execute. This can be further enriched by a *Delegate* statement, which passes responsibility to another object for executing and returning to the client when complete.
- A language built around objects that lacks inheritance is sometimes referred to as *object-based* but inheritance can be added to the features above by viewing it as a way of instantiating nested blocks.

An operational semantics for such a language can be built around the following semantic objects.

The basic threads per object are keyed by *References*:

$$\text{ObjectStore} = \text{Reference} \xrightarrow{m} \text{ObjectInformation}$$

This keeps a record of the states of all the objects that exist at a given time in the execution of the program.

Each *ObjectInformation* contains the information needed to determine the current state and activity of the object:<sup>40</sup>

$$\begin{aligned} \text{ObjectInformation} &:: \text{class} : \text{Id} \\ &\quad \sigma : \text{Store} \\ &\quad \text{mode} : \text{READY} | \text{Run} | \text{Wait} \end{aligned}$$

The local *Store* of an object simply contains the current values of its variables:

$$\text{Store} = \text{Id} \xrightarrow{m} \text{Value}$$

$$\text{Value} = [\text{Reference}] | \mathbb{Z} | \mathbb{B}$$

where the set *Reference* is infinite and  $\mathbf{nil} \notin \text{Reference}$ .

Modes of objects indicate their current activity status. Objects which are READY are not currently doing anything; method calls may be made to such objects. The other modes indicate some form of activity.

$$\begin{aligned} \text{Run} &:: \text{remainder} : \text{Statement}^* \\ &\quad \text{client} : \text{Reference} \end{aligned}$$

Objects in *Run* mode are currently executing. It is important to retain the list of statements which they have yet to execute, *remainder*, (compare with the configurations of Sect. 4.1) and the reference of the object which initiated their execution, *client*, which will be awaiting the eventual return of a value (or a special token indicating there is no return value).

$$\begin{aligned} \text{Wait} &:: \quad \text{lhs} : \text{Id} \\ &\quad \text{remainder} : \text{Statement}^* \\ &\quad \text{client} : \text{Reference} \end{aligned}$$

<sup>40</sup> The texts of object classes are stored in a separate *ClassStore*, discussion of which is postponed to the consideration of the *Program* type.

Objects waiting for a value to be returned must keep track of the (local) variable to which this value should be saved (*lhs*), the list of statements to which they will resume executing (*remainder*) and the *client* by which they were originally called.

Programs are defined as a specification of objects and some initialisation.

$$\begin{aligned} \text{Program} &:: & cs &: \text{ClassStore} \\ && \text{startingclasses} &: \text{Id}^* \\ && \text{startingmethods} &: \text{Id}^* \end{aligned}$$

The *startingclasses* sequence indicates which classes within the *ClassStore* should be initialised at program commencement and *startingmethods* indicates which methods within these classes should be executed.

*ClassStore* is the global directory of all classes in the program; the *ObjectStore* is the store of dynamic information on the extant objects; the *ClassStore* holds the static information on all possible objects.

$$\text{ClassStore} = \text{Id} \xrightarrow{m} \text{ClassInformation}$$

$$\begin{aligned} \text{ClassInformation} &:: \text{variables} : \text{Id} \xrightarrow{m} \text{Type} \\ && \text{methods} : \text{Id} \xrightarrow{m} \text{MethodInfo} \end{aligned}$$

The information here defines the variables declared in the class and their types (there are no dynamic declarations in this language) and the methods available to be called in the language. More detail need not be given on *MethodInfo* but it contains parameter information and statements to be executed for each method.

Thus the main semantic relation has the type:

$$\xrightarrow{st} : \mathcal{P}((\text{ClassStore} \times \text{ObjectStore}) \times \text{ObjectStore})$$

Once the program has commenced, the *ClassStore* and *ObjectStore* maps are globally available to the semantics during execution. However, individual objects have access to only the *ClassStore* object (to enable them to call methods in other objects) and of course their own internal store.

A full definition of COOL is available on the web<sup>41</sup> but it is a part of the message of Sect. 7 that it is possible to understand many design decisions of a programming language solely from its *semantic objects*.

## 6 Abnormal Ordering

Many programming languages contain features that bring about a non-sequential order of execution of statements. The most obvious example is the **goto** statement (attacked by Dijkstra in [Dij68] and defended by Knuth in [Knu74]) but it is certainly not the sole source of difficulty: (loop) breaks, exception mechanisms and even returns from functions or procedures present similar challenges. Expressed in denotational terms, the difficulty is that the *homomorphic rule* cannot directly apply when the meaning of a construct depends on something that is not present in the construct. Put another way, the obvious idea that the

<sup>41</sup> <http://homepages.cs.ncl.ac.uk/cliff.jones/COOL-WWW-version.pdf>.

semantics of the sequential composition of two statements should be the composition of the semantics of those two statements cannot apply when the first statement appoints as its successor a statement elsewhere.

One response from operational semantics that shows rather clearly what has to happen can be seen in VDL descriptions. In early Vienna Lab operational semantics, an explicit *control tree* recorded the text that was still to be executed; abnormal sequencing was modelled by surgery on this control tree.<sup>42</sup>

Within the denotational camp, there are two rather different responses to the challenges of abnormal ordering. Most researchers (and certainly those strongly connected to Oxford) use *Continuations*. The core idea is to recover some semblance of the homomorphic rule by making the denotation of a label represent the effect of starting execution at that label. In order to develop such denotations it is necessary to pass to every semantic function a denotation that corresponds to the execution of the remainder of the program. This makes the semantics higher order than one might expect and arguably more complicated than these specific constructs require.

In contrast, VDM denotational descriptions (and the Isabelle formulations of semantics in [NK13]) effectively extend the denotations from  $\Sigma \rightarrow \Sigma$  to have ranges that can represent abnormal results. The potential messiness caused by the fact that something more complicated than functional composition is now needed for sequential composition can be hidden by *combinators*.<sup>43</sup>

Incorporating the exit ideas into SOS descriptions is something that has not been published. It would be easy to do this explicitly with extra cases for all language constructs but this would result in the heaviness visible in [ACJ72]—much lengthier than what VDM achieves with combinators. Since the latter could be read operationally, it should be possible to find a way of adding something like the combinators to SOS rules.

An axiomatic approach to jumps is proposed in [CH72], although the authors do acknowledge that jumps may be better avoided where possible and indeed most axiomatic semantic descriptions skip the topic entirely. The essential idea is adapted from earlier (operational) work by Landin [Lan65a, Lan65b], which treated jumps like procedures whose body is the sequence of statements following the label up until the end of its enclosing block. Rather than returning control to the calling context, however, it is resumed from the end of the block enclosing the label. Clint and Hoare's approach is largely the same, although they prefer to restrict the declaration of labels (and their 'bodies') to the beginning of blocks. The rules do allow for labels to be declared anywhere within the block, with some slight added complexity. However, only one label may be declared per block, and further restrictions prevent jumping into compound and conditional statements.

<sup>42</sup> It is interesting to note that [McC66] had an explicit program counter that could be seen as a hint of what had to be done with control trees when a massive language like PL/I (complete with concurrency) had to be described.

<sup>43</sup> In [Mos11], Mosses makes the interesting link between VDM's use of such combinators and Eugenio Moggi's *monads* [Mog89]. The differences between the VDM exit scheme and continuations are teased apart by proofs of equivalence in [Jon78, Jon82].



It is interesting to note that this approach bears some obvious similarities to the continuations used in denotational semantics. Although notationally very different, the idea of a label representing computation left to be performed is at the core of both ideas.<sup>44</sup> There is also a clear comparison to the configurations used in the operational semantics of Sect. 4.1 in which the *text* of the computation yet to be executed is stored.

## 7 Closing Remarks

This section mentions some current research (Sect. 7.1), related references (Sect. 7.2) and offers some general conclusions.

### 7.1 Algebraic Semantics

Work on this topic is too recent to present a full evaluation; here only some pointers and superficial comments are offered. For sequential programs, a search for “Laws of Programming” was started in [HHJ+87]; Hoare [HvS12, HMSW11] and others [Hay16, HCM+16] build on Kozen’s *Kleene algebra with tests* [Koz97] to record algebraic laws that abstract from any detailed model of concurrent programming languages. As with *Boolean algebras*, the algebraic laws normally admit more than one model: saying, for example, that the sequence operator of semicolon is associative but non-commutative does not preclude a semantics in which statements are executed right to left.

The clear advantage of recording algebraic laws about programming constructs is the same as in classical algebra: if proofs can be conducted at that level of abstraction they are likely to be much easier and more general than any attempt to reason about a model-oriented language description. A specific example is the use made in [Hay16] of an *interchange law* to justify the equivalent of the most important Rely/Guarantee parallel introduction rule. Furthermore, Hayes and colleagues have gone on to present a *Synchronous Program Algebra* that also covers synchronous event-based concurrent languages [HCM+16]. It is interesting that there are echoes here of the *program schema* research [Pat67, LPP70] that was one of the earliest avenues of programming language research.

### 7.2 Related References

Frank de Boer has provided a proof system for POOL [dB91] which he shows to be consistent and complete with respect to an operational semantics. The assertion language works on three levels and is not first order—although it is not a higher order logic in the sense that, say, HOL is. There are also some restrictions of the POOL language.

<sup>44</sup> Indeed, in de Bakker’s book *Mathematical Theory of Program Correctness*, a book showing the use of all kinds of semantics in program proof, de Bruin gives a similar axiomatic rule but notes that it is hard to see clearly the correctness of this rule or use the rule in proofs [dBDBZ80]. Instead, a denotational-style continuations semantics is presented and proofs are built around that.

Another paper by the current authors [AJ18] looks at four complete formal descriptions of ALGOL 60, making technical comparisons as well as providing a historical context for the development of the semantic styles in general and the creation of the descriptions in particular.

Although not within the scope of this paper, which focuses on programming languages, other kinds of formal language have benefited from the application of semantic methods. Hardware description languages have been treated formally to good effect: see [Gor95] and [BJQ2000] for semantics of Verilog and the collection of papers [KB12] for VHDL. Semantic descriptions have also been written for specification languages, such as CLEAR [BG80] and Z [Spi88].

### 7.3 Conclusions

A number of the most important challenges presented by programming languages to formal description are discussed in this paper.

- The challenge of associating identifiers with variable values is solved in operational and denotational semantics with a notion of state that is essentially the same in both cases. In axiomatic semantics an explicit state is apparently avoided, but the meta-variables used in assertions in essence form an implicit state.
- In axiomatic semantics, phrase structuring in programming languages, such as that used in blocks and procedures, is handled by copying text and careful name substitution to avoid clashes. In model-oriented approaches, an abstract environment associates identifiers with locations. This is once again similar in both denotational and operational semantics.
- One area in which the semantic approaches differ significantly is handling non-determinism and concurrency. In SOS, a relation is defined economically by factoring out the non-determinism in the way in which rules match configurations. In axiomatic approaches a number of options have been explored including separation logic, temporal logic and rely/guarantee. Denotational semantics requires complex refactoring of its domain spaces.
- The description of an illustrative concurrent object-oriented language indicates that it may be easiest to use an SOS approach to bring all these aspects together in a readable form.

Clearly, there are some genuine differences in the way that semantics are recorded in the main approaches but there are also some common modelling ideas that are obscured by superficial differences of presentation.

The complexity of formally recording the complete semantics of practical programming languages—larger and more feature-rich than the one demonstrated in this paper—seems unavoidable. Unfortunately, most programming languages are not even described formally *post facto*, let alone during the design process. Sadly, most programming languages are also not very good: they are hard to learn, too packed with features whose interactions prove awkward, or their behaviour is difficult to predict. One of the authors of the current paper has several times

undertaken the task of writing a formal semantics for a language which had been designed without the benefit of a formal model. The experience bears out the argument that the payoff from formality comes from its early employment. John Reynolds often made comments such as “Formality should be the midwife of languages rather than the mortician”. With more careful use of formalism at an appropriate point in the design phase, many unfortunate problems could be avoided. Although working out a formal semantics is a non-trivial task, it takes significantly less time than building a compiler and the former provides a better basis for thought experiments than the latter. Furthermore, a wider knowledge of formal semantic techniques could result in a staged approach:

- Working out and recording the *semantic domains* of a language is an extremely cost-effective way of sorting out the fundamental concepts of a language—see the discussion in Sect. 5 and note that the semantic domains for PL/I cover less than two pages of its 100 page description [BBH+74].
- Although denotational descriptions of concurrent languages are still a subject of research, SOS descriptions provide a convenient way to make sure that the more novel aspects of updating the state of a language have been properly thought out.
- Again, it might not be practical to create a complete algebraic characterisation of a language, but thinking about the question of equivalences that should hold ought yield a language that is easier to use.
- Programmers using a language have to reason about the effects of their programs—they might do this less formally than in a textbook but their reasoning is in any case dependant on rules of inference about the constructs of the language. A statement for which it is too difficult to provide such rules is an indication that the programmer’s task has been made gratuitously difficult.

**Acknowledgements.** The authors are extremely grateful to Mike Dodds, Shmuel Tyszberowicz and Ian Hayes for constructive and detailed comments on drafts of this paper. Some of the material was also presented at HaPoC-2017 in Oxford and useful comments were made by participants. Funding for the authors’ research comes from UK EPSRC both as a PhD studentship and the *Strata* Platform grant.

## References

- [Abr10] Abrial, J.-R.: The Event-B Book. Cambridge University Press, Cambridge (2010)
- [ACJ72] Allen, C.D., Chapman, D.N., Jones, C.B.: A formal definition of ALGOL 60. Technical report 12.105, IBM Laboratory Hursley, August 1972
- [Acz82] Aczel, P.: A note on program verification. Manuscript (private communication), Manchester, January 1982
- [Acz83] Aczel, P.H.G.: On an inference rule for parallel composition. Private communication (1983)
- [AGM92] Abramsky, S., Gabbay, D.M., Maibaum, S.E. (eds.) Handbook of Logic in Computer Science: Background: Computational Structures, vol. 2. Oxford University Press Inc., New York (1992)

- [AJ18] Astarte, T.K., Jones, C.B.: Formal semantics of ALGOL 60: four descriptions in their historical context. In: De Mol, L., Primiero, G. (eds.) *Reflections on Programming Systems - Historical and Philosophical Aspects*, pp. 71–141. Springer Philosophical Studies Series (2018, in press)
- [Ame89] America, P.H.M.: The practical importance of formal semantics. In: de Bakker, J.W. (ed.) *25 jaar semantiek*. CWI (1989)
- [ANS76] ANSI: Programming language PL/I. Technical report X3.53-1976, American National Standard (1976)
- [Apt81] Apt, K.R.: Ten years of Hoare’s logic: a survey—part I. *ACM Trans. Program. Lang. Syst.* **3**(4), 431–483 (1981)
- [Apt84] Apt, K.R.: Ten years of Hoare’s logic: a survey - part II: nondeterminism. *Theor. Comput. Sci.* **28**, 83–109 (1984)
- [AR92] America, P., Rutten, J.: A layered semantics for a parallel object-oriented language. *Form. Asp. Comput.* **4**(4), 376–408 (1992)
- [Ast19] Astarte, T.K.: Formalising meaning: a history of programming language semantics. Ph.D. thesis, Newcastle University (2019, forthcoming)
- [BBG+60] Backus, J.W., et al.: Report on the algorithmic language ALGOL 60. *Numerische Mathematik* **2**(1), 106–136 (1960)
- [BBH+74] Bekič, H., Bjørner, D., Henhapl, W., Jones, C.B., Lucas, P.: A formal definition of a PL/I subset. Technical report 25.139, IBM Laboratory Vienna, December 1974
- [BG80] Burstall, R.M., Goguen, J.A.: The semantics of clear, a specification language. In: Bjørner, D. (ed.) *Abstract Software Specifications*. LNCS, vol. 86, pp. 292–332. Springer, Heidelberg (1980). [https://doi.org/10.1007/3-540-10007-5\\_41](https://doi.org/10.1007/3-540-10007-5_41)
- [BIJW75] Bekič, H., Izbicki, H., Jones, C.B., Weissenböck, F.: Some experiments with using a formal language definition in compiler development. Laboratory note LN 25.3.107, IBM Laboratory Vienna, December 1975
- [BJQ2000] Bowen, J.P., Jifeng, H., Qiwen, X.: An animatable operational semantics of the Verilog hardware description language. In: *Formal Engineering Methods*, pp. 199–207. IEEE (2000)
- [BK84] Bergstra, J.A., Klop, J.W.: Process algebra for synchronous communication. *Inf. Control* **60**(1–3), 109–137 (1984)
- [BO80] Bjørner, D., Nest, O.N. (eds.): *Towards a Formal Description of Ada*. LNCS, vol. 98. Springer, Heidelberg (1980). <https://doi.org/10.1007/3-540-10283-3>
- [BO16] Brookes, S., O’Hearn, P.W.: Concurrent separation logic. *ACM SIGLOG News* **3**(3), 47–65 (2016)
- [Bur66] Burstall, R.M.: Semantics of assignment. *Mach. Intell.* **2**, 3–20 (1966)
- [BvW98] Back, R.-J., von Wright, J.: *Refinement Calculus: A Systematic Introduction*. Springer, New York (1998). <https://doi.org/10.1007/978-1-4612-1674-2>
- [CG90] Carré, B., Garnsworthy, J.: Spark—an annotated Ada subset for safety-critical programming. In: *Proceedings of the Conference on TRI-Ada 1990*, TRI-Ada 1990, pp. 392–402. ACM (1990)
- [CH72] Clint, M., Hoare, C.A.R.: Program proving: jumps and functions. *Acta Informatica* **1**(3), 214–224 (1972)
- [CHM16] Colvin, R.J., Hayes, I.J., Meinicke, L.A.: Designing a semantic model for a wide-spectrum language with concurrency. *Form. Asp. Comput.* **29**(5), 1–22 (2016)

- [CJ06] Coleman, J.W., Jones, C.B.: Guaranteeing the soundness of rely/guarantee rules. Technical report CS-TR-955, School of Computing Science, University of Newcastle, March 2006
- [Col08] Coleman, J.W.: Constructing a tractable reasoning framework upon a fine-grained structural operational semantics. Ph.D. thesis, Newcastle University, January 2008
- [Dat82] Date, C.J.: A formal definition of the relational model. *ACM SIGMOD Rec.* **13**(1), 18–29 (1982)
- [dB91] Boer, F.S.: A proof system for the language POOL. In: de Bakker, J.W., de Roever, W.P., Rozenberg, G. (eds.) *REX 1990. LNCS*, vol. 489, pp. 124–150. Springer, Heidelberg (1991). <https://doi.org/10.1007/BFb0019442>
- [dBDBZ80] de Bakker, J.W., De Bruin, A., Zucker, J.: *Mathematical Theory of Program Correctness*, vol. 980. Prentice-Hall International, London (1980)
- [Dij68] Dijkstra, E.W.: Go to statement considered harmful. *Commun. ACM* **11**(3), 147–148 (1968)
- [Dij76] Dijkstra, E.W.: *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs (1976)
- [DMN68] Dahl, O.-J., Myhrhaug, B., Nygaard, K.: *SIMULA 67 common base language*. Technical report S-2, Norwegian Computing Center, Oslo (1968)
- [Don76] Donahue, J.E.: *Complementary Definitions of Programming Language Semantics. LNCS*, vol. 42. Springer, Heidelberg (1976). <https://doi.org/10.1007/BFb0025364>
- [dR01] de Roever, W.P.: *Concurrency Verification: Introduction to Compositional and Noncompositional Methods*. Cambridge University Press, Cambridge (2001)
- [DS90] Dijkstra, E.W., Scholten, C.S.: *Predicate Calculus and Program Semantics*. Springer, New York (1990). <https://doi.org/10.1007/978-1-4612-3228-5>
- [Flo67] Floyd, R.W.: Assigning meanings to programs. In: *Proceedings of Symposium in Applied Mathematics. Mathematical Aspects of Computer Science*, vol. 19, pp. 19–32. American Mathematical Society (1967)
- [Gor75] Gordon, M.: *Operational reasoning and denotational semantics*. Technical report STAN-CS-75-506, Computer Science Department, Stanford University, August 1975
- [Gor95] Gordon, M.: The semantic challenge of Verilog HDL. In: *Proceedings of the Tenth Annual IEEE Symposium on Logic in Computer Science, LICS 1995*, pp. 136–145. IEEE (1995)
- [GP99] Gabbay, M., Pitts, A.: A new approach to abstract syntax involving binders. In: *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science, LICS 1999*. IEEE Computer Society (1999)
- [GR83] Goldberg, A., Robson, D.: *Smalltalk-80: The Language and Its Implementation*. Addison-Wesley, Boston (1983)
- [Hay16] Hayes, I.J.: Generalised rely-guarantee concurrency: an algebraic foundation. *Form. Asp. Comput.* **28**(6), 1057–1078 (2016)
- [HCM+16] Hayes, I.J., Colvin, R.J., Meinicke, L.A., Winter, K., Velykis, A.: An algebra of synchronous atomic steps. In: Fitzgerald, J., Heitmeyer, C., Gnesi, S., Philippou, A. (eds.) *FM 2016. LNCS*, vol. 9995, pp. 352–369. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-48989-6\\_22](https://doi.org/10.1007/978-3-319-48989-6_22)
- [HHJ+87] Hoare, C.A.R., et al.: Laws of programming. *Commun. ACM* **30**(8), 672–687 (1987). See Corrigenda in *Commun. ACM* **30**(9), 770

- [HJ73] Hanford, K.V., Jones, C.B.: Dynamic syntax: a concept for the definition of the syntax of programming languages. In: *Annual Review in Automatic Programming*, vol. 7, pp. 115–140. Pergamon (1973)
- [HJ08] Hughes, J.R.D., Jones, C.B.: Reasoning about programs via operational semantics: requirements for a support system. *Autom. Softw. Eng.* **15**(3–4), 299–312 (2008)
- [HJC14] Hayes, I.J., Jones, C.B., Colvin, R.J.: Laws and semantics for rely-guarantee refinement. Technical report CS-TR-1425, Newcastle University, July 2014
- [HMRC87] Holt, R.C., Matthews, P.A., Rosselet, J.A., Cordy, J.R.: *The Turing Programming Language: Design and Definition*. Prentice-Hall Inc., Upper Saddle River (1987)
- [HMSW11] Hoare, C.A.R., Möller, B., Struth, G., Wehrman, I.: Concurrent Kleene Algebra and its foundations. *J. Log. Algebr. Program.* **80**(6), 266–296 (2011)
- [HMT87] Harper, R., Milner, R., Tofte, M.: The semantics of standard ML: version 1, Laboratory for Foundations of Computer Science, Department of Computer Science, University of Edinburgh (1987). Hard copy
- [Hoa69] Hoare, C.A.R.: An axiomatic basis for computer programming. *Commun. ACM* **12**(10), 576–580 (1969)
- [Hoa71a] Hoare, C.A.R.: Procedures and parameters: an axiomatic approach. In: Engeler, E. (ed.) *Symposium on Semantics of Algorithmic Languages*. LNM, vol. 188, pp. 102–116. Springer, Berlin (1971)
- [Hoa71b] Hoare, C.A.R.: Proof of a program: FIND. *Commun. ACM* **14**(1), 39–45 (1971)
- [Hoa72a] Hoare, C.A.R.: A note on the FOR statement. *BIT* **12**(3), 334–341 (1972)
- [Hoa72b] C.A.R. Hoare. Towards a theory of parallel programming. In *Operating System Techniques*, pages 61–71. Academic Press, 1972
- [Hoa73] Hoare, C.A.R.: Hints on programming language design. Invited Address at SIGACT/SIGPLAN Symposium on Principles of Programming Languages, Boston, October 1973
- [Hoa85] Hoare, C.A.R.: *Communicating Sequential Processes*. Prentice-Hall, Upper Saddle River (1985)
- [Hug11] Hughes, J.R.D.: Reasoning about programs using operational semantics and the role of a proof support tool. Ph.D. thesis, Newcastle University (2011)
- [HvS12] Hoare, T., van Staden, S.: In praise of algebra. *Form. Asp. Comput.* **24**(4–6), 423–431 (2012)
- [INM88] INMOS. *occam 2: Reference Manual*. Prentice Hall (1988)
- [Izb75] Izicki, H.: On a consistency proof of a chapter of a formal definition of a PL/I subset. Technical report TR 25.142, IBM Laboratory Vienna, February 1975
- [JA16] Jones, C.B., Astarte, T.K.: An exegesis of four formal descriptions of ALGOL 60. Technical report CS-TR-1498 School of Computer Science, Newcastle University, September 2016. Forthcoming as a paper in the HaPoP 2016 Proceedings
- [JH16] Jones, C.B., Hayes, I.J.: Possible values: exploring a concept for concurrency. *J. Log. Algebraic Methods Program.* **85**, 972–984 (2016)
- [JHC15] Jones, C.B., Hayes, I.J., Colvin, R.J.: Balancing expressiveness in formal approaches to concurrency. *Form. Asp. Comput.* **27**(3), 465–497 (2015)

- [JL71] Jones, C.B., Lucas, P.: Proving correctness of implementation techniques. In: Engeler, E. (ed.) *Symposium on Semantics of Algorithmic Languages*. LNM, vol. 188, pp. 178–211. Springer, Heidelberg (1971). <https://doi.org/10.1007/BFb0059698>
- [JLRW05] Jones, C.B., Lomet, D., Romanovsky, A., Weikum, G.: The atomic manifesto: a story in four quarks. *ACM SIGMOD Rec.* **34**(1), 63–69 (2005)
- [Jon69] Jones, C.B.: A proof of the correctness of some optimising techniques. Technical report LN 25.3.051, IBM Laboratory, Vienna, June 1969
- [Jon76] Jones, C.B.: Formal definition in compiler development. Technical report 25.145, IBM Laboratory Vienna, February 1976
- [Jon78] Jones, C.B.: Denotational semantics of goto: an exit formulation and its relation to continuations. In Bjørner and Jones [BJ78], pp. 278–304
- [Jon80] Jones, C.B.: *Software Development: a Rigorous Approach*. Prentice Hall International, Englewood Cliffs (1980)
- [Jon81] Jones, C.B.: Development methods for computer programs including a notion of interference. Ph.D. thesis, Oxford University, June 1981. Printed as: Programming Research Group, Technical Monograph 25
- [Jon82] Jones, C.B.: More on exception mechanisms. In: Bjørner and Jones [BJ82], Chap. 5, pp. 125–140
- [Jon90] Jones, C.B.: *Systematic Software Development using VDM*, 2nd edn. Prentice Hall International, Upper Saddle River (1990)
- [Jon03] Jones, C.B.: The early search for tractable ways of reasoning about programs. *IEEE Ann. Hist. Comput.* **25**(2), 26–49 (2003)
- [JP11] Jones, C.B., Pierce, K.G.: Elucidating concurrent algorithms via layers of abstraction and reification. *Form. Asp. Comput.* **23**(3), 289–306 (2011)
- [JY15] Jones, C.B., Yatapanage, N.: Reasoning about separation using abstraction and reification. In: Calinescu, R., Rumpe, B. (eds.) *SEFM 2015*. LNCS, vol. 9276, pp. 3–19. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-22969-0\\_1](https://doi.org/10.1007/978-3-319-22969-0_1)
- [Kah87] Kahn, G.: Natural semantics. In: Brandenburg, F.J., Vidal-Naquet, G., Wirsing, M. (eds.) *STACS 1987*. LNCS, vol. 247, pp. 22–39. Springer, Heidelberg (1987). <https://doi.org/10.1007/BFb0039592>
- [KB12] Kloos, C.D., Breuer, P.: *Formal Semantics for VHDL*. Springer, Heidelberg (2012). <https://doi.org/10.1007/978-1-4615-2237-9>
- [Kin69] King, J.C.: A program verifier. Ph.D. thesis, Department of Computer Science, Carnegie-Mellon University (1969)
- [Knu64] Knuth, D.E.: Man or boy. *ALGOL Bull.* **17**(7) (1964)
- [Knu68] Knuth, D.E.: Semantics of context-free languages. *Theory Comput. Syst.* **2**(2), 127–145 (1968)
- [Knu74] Knuth, D.E.: Structured programming with GO TO statements. Technical report STAN-CS-74-416, Computer Science Dept, Stanford University, May 1974
- [Koz97] Kozen, D.: Kleene algebra with tests. *ACM Trans. Program. Lang. Syst.* **19**(3), 427–443 (1997)
- [Lab66] Vienna Laboratory: Formal definition of PL/I (Universal Language Document No. 3). Technical report 25.071, IBM Laboratory Vienna, December 1966
- [Lan65a] Landin, P.J.: A correspondence between ALGOL 60 and Church’s lambda-notation: part I. *Commun. ACM* **8**(2), 89–101 (1965)
- [Lan65b] Landin, P.J.: A correspondence between ALGOL 60 and Church’s lambda-notation: part II. *Commun. ACM* **8**(3), 158–167 (1965)

- [Lau71] Lauer, P.E.: Consistent formal theories of the semantics of programming languages. Ph.D. thesis, Queen's University of Belfast (1971). Printed as TR 25.121, IBM Lab. Vienna
- [LPP70] Luckham, D.C., Park, D.M.R., Paterson, M.S.: On formalised computer programs. *J. Comput. Syst. Sci.* **4**(3), 220–249 (1970)
- [Luc68] Lucas, P.: Two constructive realisations of the block concept and their equivalence. Technical report TR 25.085, IBM Laboratory Vienna, June 1968
- [LW69] Lucas, P., Walk, K.: On the formal description of PL/I. *Annu. Rev. Autom. Program.* **6**, 105–182 (1969)
- [McC63] McCarthy, J.: Towards a mathematical science of computation. In: IFIP Congress, vol. 62, pp. 21–28 (1962)
- [McC66] McCarthy, J.: A formal description of a subset of ALGOL. In: *Formal Language Description Languages for Computer Programming*, pp. 1–12. North-Holland (1966)
- [Men64] Mendelson, E.: *Introduction to Mathematical Logic*. van Norstrand (1964)
- [Mey88] Meyer, B.: *Object-Oriented Software Construction*. Prentice-Hall, Upper Saddle River (1988)
- [Mil89] Milner, R.: *Communication and Concurrency*. Prentice Hall, Upper Saddle River (1989)
- [Mog89] Moggi, E.: An abstract view of programming languages. Ph.D. thesis, Laboratory for the Foundation of Computer Science, Edinburgh University (1989)
- [Mor94] Morgan, C.C.: *Programming from Specifications*, 2nd edn. Prentice Hall, Upper Saddle River (1994)
- [Mos11] Mosses, P.D.: VDM semantics of programming languages: combinators and monads. *Form. Asp. Comput.* **23**(2), 221–238 (2011)
- [MP67] McCarthy, J., Painter, J.: Correctness of a compiler for arithmetic expressions. *Math. Asp. Comput. Sci.* **19** (1967)
- [MS74] Milne, R., Strachey, C.: A theory of programming language semantics. Privately circulated (1974). Submitted for the Adams Prize
- [MS76] Milne, R., Strachey, C.: *A Theory of Programming Language Semantics (Parts A and B)*. Chapman and Hall, Boca Raton (1976)
- [NK13] Nipkow, T., Klein, G.: *Concrete Semantics. A Proof Assistant Approach*. Springer, Cham (2013)
- [NN92] Nielson, H.R., Nielson, F.: *Semantics with Applications: A Formal Introduction*. Wiley, New York (1992)
- [O'H07] O'Hearn, P.W.: Resources, concurrency and local reasoning. *Theor. Comput. Sci.* **375**(1–3), 271–307 (2007)
- [Pag81] Pagan, F.G.: *Formal Specification of Programming Languages*. Prentice-Hall, Upper Saddle River (1981)
- [Pai67] Painter, J.A.: Semantic correctness of a compiler for an ALGOL-like language. Technical report AI Memo 44, Computer Science Department, Stanford University, March 1967
- [Par10] Parkinson, M.: The next 700 separation logics. In: Leavens, G.T., O'Hearn, P., Rajamani, S.K. (eds.) *VSTTE 2010. LNCS*, vol. 6217, pp. 169–182. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-15057-9\\_12](https://doi.org/10.1007/978-3-642-15057-9_12)
- [Pat67] Paterson, M.S.: Equivalence problems in a model of computation. Ph.D. thesis, University of Cambridge (1967)



- [Pie02] Pierce, B.C.: Types and Programming Languages. MIT Press, Cambridge (2002)
- [Plo76] Plotkin, G.D.: A powerdomain construction. *SIAM J. Comput.* **5**, 452–487 (1976)
- [Plo81] Plotkin, G.D.: A structural approach to operational semantics. Technical report DAIMI FN-19, Aarhus University (1981)
- [Pra65] Prawitz, D.: Natural Deduction: A Proof-Theoretical Study. Dover Publications, New York (1965)
- [Rey78] Reynolds, J.C.: Syntactic control of interference. In: Proceedings of Fifth POPL, pp. 39–46. ACM (1978)
- [Rey89] Reynolds, J.C.: Syntactic control of interference part 2. In: Ausiello, G., Dezani-Ciancaglini, M., Della Rocca, S.R. (eds.) ICALP 1989. LNCS, vol. 372, pp. 704–722. Springer, Heidelberg (1989). <https://doi.org/10.1007/BFb0035793>
- [RR62] Randell, B., Russell, L.J.: Discussions on ALGOL translation at Mathematisch Centrum. English Electric Report W/AT, 841 (1962)
- [San99] Sangiorgi, D.: Typed pi-calculus at work: a correctness proof of Jones’s parallelisation transformation on concurrent objects. *TAPOS* **5**(1), 25–33 (1999)
- [Sco80] Scott, D.: Lambda calculus: some models, some philosophy. *Stud. Log. Found. Math.* **101**, 223–265 (1980)
- [Smy76] Smyth, M.B.: Powerdomains. Technical report, Department of Computer Science, University of Warwick, May 1976
- [Spi88] Spivey, J.M.: Understanding Z—A Specification Language and its Formal Semantics. Cambridge Tracts in Computer Science 3. Cambridge University Press (1988)
- [Ste66] Steel, T.B.: Formal Language Description Languages for Computer Programming. North-Holland, London (1966)
- [Sto77] Stoy, J.E.: Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory. MIT Press, Cambridge (1977)
- [SW01] Sangiorgi, D., Walker, D.: The  $\pi$ -Calculus: A Theory of Mobile Processes. Cambridge University Press, Cambridge (2001)
- [Tur49] Turing, A.M.: Checking a large routine. In: Report of a Conference on High Speed Automatic Calculating Machines, pp. 67–69. University Mathematical Laboratory, Cambridge, June 1949
- [Tur09] Turner, R.: The meaning of programming languages. *Am. Philos. Assoc. Newsl. Philos. Comput.* **9**(1), 2–6 (2009)
- [vdH17] van den Hove, G.: New insights from old programs: the structure of the first ALGOL 60 system. Ph.D. thesis, University of Amsterdam (2017)
- [vWMPK69] van Wijngaarden, A., Mailloux, B.J., Peck, J.E.L., Koster, C.H.A.: Report on the algorithmic language ALGOL 68. Mathematisch Centrum, Amsterdam, October 1969. Second printing, MR 101
- [WAB+68] Walk, K., et al.: Abstract syntax and interpretation of PL/I. Technical report 25.082, IBM Laboratory Vienna, ULD Version II, June 1968
- [Wal67] Walk, K.: Minutes of the 1st meeting of IFIP WG 2.2 on Formal Language Description Languages. Kept in the van Wijngaarden archive: Held in Porto Conte, Alghero, Sardinia (1967)
- [Wal69] Walk, K.: Minutes of the 3rd Meeting of IFIP WG 2.2 on Formal Language Description Languages, April 1969. Held in Vienna, Austria
- [Wei75] Weissenböck, F.: A formal interface specification. Technical report TR 25.141, IBM Laboratory Vienna, February 1975

- [Win93] Winskel, G.: The Formal Semantics of Programming Languages. The MIT Press (1993). ISBN 0-262-23169-7
- [Woo93] Woodman, M.: A taste of the Modula-2 standard. ACM SIGPLAN Not. **28**(9), 15–24 (1993)
- [Yon90] Yonezawa, A. (ed.): ABCL: An Object-Oriented Concurrent System. MIT Press, Cambridge (1990). ISBN 0-262-24029-7