



Fifty years of Hoare’s logic

Krzysztof R. Apt^{1,2} and Ernst-Rüdiger Olderog³

¹CWI, Amsterdam, The Netherlands

²MIMUW, University of Warsaw, Warsaw, Poland

³University of Oldenburg, Oldenburg, Germany

Contents

1	Introduction	752
2	Precursors	753
2.1	Turing	753
2.2	Floyd	754
3	Hoare’s Contributions	755
3.1	Reasoning about while programs	755
3.2	Reasoning about recursive procedures	757
3.3	Reasoning about termination	759
4	Soundness and Completeness Matters	761
4.1	Preliminaries	761
4.2	Soundness	761
4.3	Completeness	762
5	Fine-tuning the Approach	764
5.1	Adaptation rules	764
5.2	Subscripted and local variables	766
5.3	Parameter mechanisms and procedure calls	769
6	Reasoning about Arbitrary Procedures	770
6.1	Completeness results for recursive procedures	770
6.2	Clarke’s incompleteness result	772
6.3	Clarke’s language L_4	773
6.4	The characterization problem	774
7	Nondeterministic and Probabilistic Programs	775
7.1	Reasoning about nondeterminism	775
7.2	Reasoning about fairness	777
7.3	Probabilistic programs	778
8	Parallel and Distributed Programs	779
8.1	Reasoning about parallel programs	779
8.2	Reasoning about distributed programs	783
9	Object-oriented Programs	786
9.1	Language characteristics	786
9.2	Reasoning about object-oriented programs	788
9.3	Advanced topics in the verification of object-oriented programs	789

10	Alternative Approaches	790
10.1	Weakest precondition semantics and systematic program development	790
10.2	Specifying in Hoare's logic	792
10.3	Programming from specifications	793
10.4	Algorithmic logic and dynamic logic	794
10.5	Temporal logic and model checking	795
10.6	Separation logic	796
10.7	Relational Hoare logic	798
11	Final Remarks: a Summary and an Assessment	799
A	Turing's example	801

Abstract. We present a history of Hoare's logic.

1. Introduction

Hoare's logic is a formalism allowing us to reason about program correctness. It was introduced fifty years ago in the seminal article [Hoa69] of Tony Hoare that focused on a small class of **while** programs, and was soon extended by him in [Hoa71a] to programs allowing local variables and recursive procedures. This approach became the most influential method of verifying programs, mainly because its syntax-oriented style made it possible to extend it to almost any type of programs. Also, thanks to parallel developments in program semantics, this approach leans itself naturally to a rigorous analysis based on the methods of mathematical logic. Since then, several books appeared that discuss Hoare's logic, or at least have a chapter on it: [dB80, LS87, TZ88, Fra92, Win93, AFPdS11, Ben12], to name a few.

More than thirty years ago two surveys of Hoare's logic appeared, [Apt81], concerned with deterministic programs, and [Apt84], concerned with nondeterministic programs. At the beginning of the nineties an extensive survey [Cou90] was published that also included an account of verification of parallel programs and a discussion of other approaches to program verification.

A systematic exposition of Hoare's logics for deterministic, nondeterministic and parallel programs appeared in our book [AO91]. The last edition of it, [AdBO09], written jointly with F.S. de Boer, extended the presentation to recursive procedures and object-oriented programs. In this paper we occasionally rely on the material presented in this book, notably to structure the presentation.

Given that the literature on the subject is really vast, we would like to clarify the scope of this survey. Our objective is to trace the main developments in Hoare's logic that originated with the original article [Hoa69] and resulted in proof systems dealing with several classes of programs. To realize this natural research programme various difficulties, notably concerning the logical status of the proposed proof systems, had to be resolved. We clarify here these and other complications and explain how they were addressed.

In particular, we discuss the problem of completeness, the issues concerning local variables, parameter mechanisms, and auxiliary rules. We also discuss various alternatives to Hoare's logic that were proposed in the literature. Given the scope of this exposition we discuss only briefly (in the final section) the use of automated reasoning in the context of Hoare's logic. This important sequel to the research reported here deserves a separate survey.

According to Google Scholar, the original article [Hoa69] has been cited more than 7000 times. This forced us to make some selection in the presented material. Some omissions, such as the treatment of the nowadays less often used **goto** statement or coroutines, were dictated by our effort to trace and explain the developments that withstood the test of time.

Further, we do not introduce any program semantics. Consequently, we do not prove any soundness or completeness results. Instead, we focus on a systematic account of the established results combined with an explanation of the reasons why some concepts were introduced, and on a discussion of some, occasionally subtle, ways Hoare's logic differs from customary logics.




Fig. 1. Turing's flowchart, reconstructed by F.L. Morris and C.B. Jones [MJ84]

We begin the exposition by discussing in the next section the contributions to program verification by Alan Turing and Robert Floyd that preceded those of Hoare. Then, in Sect. 3, we discuss Hoare's initial contributions that focused on the **while** programs and programs with recursive procedures, though we extend the exposition by an account of program termination. Next, we discuss in Sect. 4 the soundness and completeness of the discussed proof systems. An essential difference between Hoare's logic and first-order logic has to do with the features specific to programming languages, such as subscripted variables, local variables, and parameter mechanisms. We discuss these matters in Sect. 5. This provides a natural starting point for an account of verification of programs with arbitrary procedures, notably procedures that allow procedures as parameters. This forms the subject of Sect. 6.

In Sect. 7 we discuss verification of nondeterministic programs, the corresponding issue of fairness, and verification of probabilistic programs. Then, in Sect. 8 we focus on the verification of parallel and distributed programs. Next, in Sect. 9, we provide an account of verification of object-oriented programs. In the final two Sects. 10 and 11, we shed light on alternative approaches to program verification, discuss briefly the use of automated verification in relation to Hoare's logic, and attempt to explain and assess the impact of Hoare's logic.

2. Precursors

2.1. Turing

The concern about correctness of computer programs is as old as computers themselves. Tracing the history of program verification is not the focus of this survey, so we only discuss here two key contributions preceding those of Tony Hoare. Readers interested in the early history of this subject are referred to the paper [Jon03] by C.B. Jones.

In 1949, Alan Turing gave a presentation entitled "Checking a Large Routine" at a conference in Cambridge U.K. at the occasion of the launching of the computer EDSAC (Electronic Delay Storage Automatic Calculator), published as [Tur49]. F.L. Morris and C.B. Jones recovered [MJ84] the original typescript of Turing's presentation and made it available for a wider audience, thereby correcting several typing errors.

Turing started by asking

"How can one check a routine in the sense of making sure that it is right?"

and proposed that

"... the programmer should make a number of definite assertions which can be checked individually, and from which the correctness of the whole programme easily follows."

Turing demonstrated his ideas for a flowchart program with nested loops computing the factorial $n!$ of a given natural number n , where multiplication is achieved by repeated addition; see Fig. 1. Note that the effect of a command in the flowchart is represented by an equation like $u' = u + v$, where the primed variable u' denotes the value of the variable u after the execution of the command.

STORAGE LOCATION	(INITIAL) Ⓐ k = 6	Ⓑ k = 5	Ⓒ k = 4	(STOP) Ⓓ k = 0	Ⓔ k = 3	Ⓕ k = 1	Ⓖ k = 2
s		r	r		s	s + 1	s
r		n	n		r	r	r
n	n	⌊	⌊	n	n	n	n
u		⌊	⌊	⌊n	s⌊	(s + 1)⌊	(s + 1)⌊
v		⌊	⌊		⌊	⌊	⌊
27							
28							
29							
30							
31							
	TO Ⓑ WITH r' = 1 u' = 1	TO Ⓒ	TO Ⓓ IF r = n TO Ⓔ IF r < n		TO Ⓕ	TO Ⓑ WITH r' = r + 1 IF s ≥ r TO Ⓔ WITH s' = s + 1 IF s < r	TO Ⓖ

Fig. 2. Turing’s assertions, reconstructed by F.L. Morris and C.B. Jones [MJ84]. Turing writes $\lfloor n$ for the factorial of n

Today, this notation is still in use in logical representations of computation steps like in the specification language Z (see, e.g., [Spi92]) and bounded model checking.

Turing referred already to *assertions*. In the example he presented them in the form of a table referring to the numbers of the locations storing the variables s, r, n, u, v , see Fig. 2. From today’s viewpoint these assertions are admittedly very specific and difficult to read.

Turing was not only concerned with delivering correct values, but also with termination. He wrote

“Finally the checker has to verify that the process comes to an end. Here again he should be assisted by the programmer giving a further definite assertion to be verified. This may take the form of a quantity which is asserted to decrease continually and vanish when the machine stops.”

This refers already to the concept of a termination function. Turing stated a global termination function for the example program, i.e., an integer expression yielding a non-negative value that decreases with every step of the program.

Summarizing, Turing introduced the notions of assertions and termination functions, but did not state loop invariants and local termination functions for the two loops of the program. Still, as we show in the Appendix, his approach can be represented within the framework of Hoare’s logic.

2.2. Floyd

Robert Floyd was the first to propose a fully formal method for proving the correctness of flowchart programs known as *inductive assertions method* in [Flo67]. (Floyd actually mentions that his work is “based on ideas of Perlis and Gorn, and may have made their earliest appearance in an unpublished paper by Gorn.”) Here the assertions are logical formulas in terms of the variables appearing in the flowcharts. The begin of the flowchart is annotated with an assertion stating the assumptions under which the flowchart is supposed to work. The end of the flowchart is annotated with an assertion specifying the desired result. To verify that these input-output annotations are correct, each loop of the flowchart needs to be annotated at some selected point with an assertion that should hold whenever the control reaches this point. The assertion should thus be an *invariant* at this selected point. Floyd states rules how to verify this by completing the flowchart so that there is at least one assertion between any two subsequent statements. The rules explain how to modify a given assertion when passing a test statement and when passing an assignment statement. When two assertions are adjacent to the same arc then the logical implication has to hold in the direction of the arc.

In Fig. 3 we show Turing’s example as a flowchart with annotations according to Floyd’s method. At the begin B of the flowchart the annotation $n \geq 1$ states the assumption of the computation, at the end E the annotation $v = n!$ specifies the desired result that should hold whenever the computation reaches E . To verify that this annotation is correct, in every loop a point needs to be selected and annotated with an *invariant*. In this example, we select the bullet points and annotate them with the assertions

$$P_1 \equiv v = r! \wedge u = r! \wedge 1 \leq r \leq n$$

and

$$P_2 \equiv v = r! \wedge u = s \cdot v \wedge 1 \leq s \leq r + 1 \leq n.$$

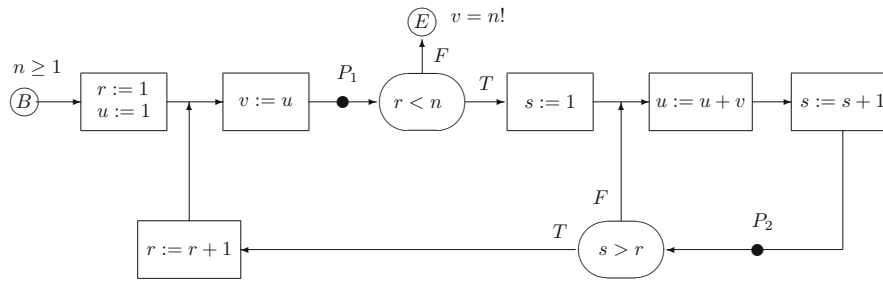


Fig. 3. Turing's example as a flowchart with annotations according to Floyd's method, where P_1 and P_2 are the following invariants: $P_1 \equiv v = r! \wedge u = r! \wedge 1 \leq r \leq n$ and $P_2 \equiv v = r! \wedge u = s \cdot v \wedge 1 \leq s \leq r + 1 \leq n$. See also [dB75]

While the inductive assertions method is a natural approach to program verification, it is limited to programs that can be represented as flowcharts. This obvious deficiency was overcome by Hoare by proposing an approach that is applicable to programs written in the customary textual form.

3. Hoare's Contributions

3.1. Reasoning about while programs

To reason about programs Hoare introduced in [Hoa69] a new notation

$$P \{S\} Q,$$

with the interpretation

"If the assertion P is true before initiation of a program S , then the assertion Q will be true on its completion."

Nowadays one rather writes

$$\{P\} S \{Q\}$$

so that additional assertions can be freely inserted in the program text, by putting the $\{\cdot\}$ brackets around them. Such a possibility will turn out to be especially important when reasoning about parallel programs. In what follows we shall use the latter notation. In this context P is referred to as a *precondition* and Q as a *postcondition*.

Subsequently Hoare introduced an axiom to reason about the assignment statement and the proof rules to reason about the program composition and the **while** statement. He also introduced two consequence rules, now combined into one, that allow one to strengthen the precondition and to weaken the postcondition. He then used these axioms and rules to establish correctness of the following simple program, let us call it DIV, that finds "the quotient q and a remainder r obtained on dividing x by y ":

$$r := x; \quad q := 0; \quad \mathbf{while} \ y \leq r \ \mathbf{do} \ r := r - y; \quad q := 1 + q \ \mathbf{od}.$$

All variables are assumed to range over the nonnegative integers.

In what follows we review these steps. We view assertions as formulas in some first-order language and assume from the reader some rudimentary knowledge of basic concepts of logic. The assignment axiom has the form:

ASSIGNMENT

$$\{P[x := t]\} x := t \{P\},$$

where x is a variable, t is an expression, and $P[x := t]$ is the result of substituting t for all free occurrences of x in P .

Line number	Formal proof	Justification
1	$\mathbf{true} \rightarrow x = x + y \cdot 0$	logic
2	$\{x = x + y \cdot 0\} r := x \{x = r + y \cdot 0\}$	ASSIGNMENT
3	$\{x = r + y \cdot 0\} q := 0 \{x = r + y \cdot q\}$	ASSIGNMENT
4	$\{\mathbf{true}\} r := x \{x = r + y \cdot 0\}$	CONSEQUENCE (1,2)
5	$\{\mathbf{true}\} r := x; q := 0 \{x = r + y \cdot q\}$	COMPOSITION (4,3)
6	$x = r + y \cdot q \wedge y \leq r \rightarrow x = (r - y) + y \cdot (1 + q)$	logic
7	$\{x = (r - y) + y \cdot (1 + q)\} r := r - y \{x = r + y \cdot (1 + q)\}$	ASSIGNMENT
8	$\{x = r + y \cdot (1 + q)\} q := 1 + q \{x = r + y \cdot q\}$	ASSIGNMENT
9	$\{x = (r - y) + y \cdot (1 + q)\} r := r - y; q := 1 + q \{x = r + y \cdot q\}$	COMPOSITION (7,8)
10	$\{x = r + y \cdot q \wedge y \leq r\} r := r - y; q := 1 + q \{x = r + y \cdot q\}$	CONSEQUENCE (6,9)
11	$\{x = r + y \cdot q\} \mathbf{while} \ y \leq r \ \mathbf{do} \ r := r - y; q := 1 + q \ \mathbf{od} \ \{\neg y \leq r \wedge x = r + y \cdot q\}$	WHILE (10)
12	$\{\mathbf{true}\} r := x; q := 0; \mathbf{while} \ y \leq r \ \mathbf{do} \ r := r - y; q := 1 + q \ \mathbf{od} \ \{\neg y \leq r \wedge x = r + y \cdot q\}$	COMPOSITION (5,11)

The arguments in the right column of the rules refer to the line numbers to which they were applied and ‘logic’ indicates that the relevant formulas are true (Hoare referred to specific axioms of Peano arithmetic).

Fig. 4. Correctness proof of the DIV program

The already mentioned consequence rule has the following form:

CONSEQUENCE

$$\frac{P \rightarrow P_1, \{P_1\} S \{Q_1\}, Q_1 \rightarrow Q}{\{P\} S \{Q\}}$$

Here it is assumed that the mentioned implications can be established in some further unspecified proof system exclusively concerned with the assertions. (Hoare just referred to \rightarrow as ‘logical implication’.)

The final two rules were:

COMPOSITION

$$\frac{\{P\} S_1 \{R\}, \{R\} S_2 \{Q\}}{\{P\} S_1; S_2 \{Q\}}$$

and

WHILE

$$\frac{\{P \wedge B\} S \{P\}}{\{P\} \mathbf{while} \ B \ \mathbf{do} \ S \ \mathbf{od} \ \{P \wedge \neg B\}}$$

Nowadays one refers to the assertion P that satisfies the premise of this rule as a *loop invariant*.

Hoare’s correctness proof of the DIV program is presented in Fig. 4. (Hoare wrote the postcondition of the conclusion of the WHILE rule as $\neg B \wedge P$ and this is how it is recorded in Fig. 4.) It yields the desired conclusion that q is the quotient and r the remainder resulting from dividing x by y . The crucial step in this proof is line 10 that clarifies the role played by the assertion $x = r + y \cdot q$. This line establishes that $x = r + y \cdot q$ is a loop invariant of the considered **while** statement and its discovery is essential for the proof to succeed.

As pointed out in [JR10] the assignment axiom was originally proposed in [Kin69], the PhD thesis of J. King. From [Flo67] one can distill a more complex assignment axiom

ASSIGNMENT I

$$\{P\} x := t \ (\exists y : (P[x := y] \wedge x = t[x := y])).$$

that reasons “forward” starting from the precondition P .

The striking simplicity of the ASSIGNMENT axiom reveals a close relation between the assignment statement and the substitution operation. This is achieved, in contrast to Floyd’s approach, by reasoning ‘backwards’, so starting from the postcondition P . The adoption of this axiom by Hoare probably influenced a couple of years

later Edsger W. Dijkstra to propose the weakest precondition semantics that adopted this reasoning ‘backward’ to all program statements. We shall discuss this alternative approach to program verification in Sect. 10. From the mathematical point of view, Hoare’s proof rules and axioms form an unusual mix: the assignment axiom adopts ‘backward’ reasoning, while the proof rules embrace ‘forward’ reasoning, in the sense that they maintain the view that a program is a transformation of a precondition to a postcondition.

Historically, Hoare’s approach was preceded by P. Naur’s paper [Nau66] in which a simple ALGOL 60 program that finds the greatest element in an array of numbers was verified using ‘snapshots’, which are in fact assertions. However, Naur’s reasoning is not backed up by any proof rules. His paper closes with this remark: “Similar concepts have been developed independently by Robert W. Floyd (unpublished paper, communicated privately).” that clearly refers to Floyd’s paper discussed in the previous subsection.

Hoare’s paper turned out to be the beginning of a far reaching change in reasoning about programs, resulting from moving from flowcharts to programs expressed in the customary textual form. This opened the way to reasoning about programs that cannot be readily expressed as flowcharts, for example, recursive procedures or programs with variable declarations. Also it made it possible to adopt a syntax-directed reasoning about programs by using their structure as guidance in organizing the proof.

A related, implicit, feature of the proof system proposed by Hoare is that it encourages program development by allowing one to first specify the desired preconditions and postconditions of a program component and subsequently to look for a program fragment for which the corresponding correctness statement can be established. Hoare took a lead in this novel view of program correctness by publishing in [Hoa71b] a correctness proof of the FIND program, the purpose of which is to find the f th largest element of an array $A[1 : N]$ by rearranging its elements so that upon termination

$$A[1], A[2], \dots, A[f - 1] \leq A[f] \leq A[f + 1], \dots, A[N].$$

The program is very subtle—it uses a triply nested **while** loop—and as a result its correctness proof is highly nontrivial. The proof is not carried out in the proof system of [Hoa69] but from the way it is written it is clear that it can be done so. In fact, Hoare refers in a number of places to *invariants* that he defines as formulas that remain true throughout the execution of the program independently of the values of the program variables.

A similar in style contribution is [Hoa72b], in which a correctness proof was given of a program encoding the sieve of Eratosthenes. The difference was that the program was developed together with its correctness proof and presented using non-recursive procedures and classes, drawing on the contemporary works of E.W. Dijkstra on structured programming and O.J. Dahl on the object-oriented programming language SIMULA 67, which appeared as chapters in [DDH72]. These two contributions of Hoare, [Hoa71b] and [Hoa72b], showed that his original logic could be seen not only as a tool to verify programs but also as a guide to design correct programs. These ideas were further developed by Dijkstra, notably in his book [Dij76a].

3.2. Reasoning about recursive procedures

Let us continue with another milestone in the history of Hoare’s logic. Foley and Hoare established correctness of the program Quicksort [FH71], originally proposed by Hoare in [Hoa61]. Foley and Hoare stated:

“The purpose of the program Quicksort is to sort the elements $a[m]$ to $a[n]$ of an array into ascending order, while leaving untouched those below $a[m]$ and above $a[n]$.”

The main difficulty was that Quicksort uses recursion. (Actually it was the first non-trivial example of a successful use of recursion in imperative programming.) This required appropriate proof rules that were introduced by Hoare in [Hoa71a].

In what follows, given a program S we denote by $change(S)$ the set of variables that are subject to change in it. In general, this set depends on the initial state of the program, but one can define syntactically a superset of it that takes care correctly of local variables and multiple procedure declarations. In what follows we mean by $change(S)$ this syntactically defined superset.

Further, we use $\mathbf{proc} p(\mathbf{x} : \mathbf{v}) : S$ to denote the declaration of a procedure p with the body S and two disjoint lists of distinct formal parameters: \mathbf{x} is the list of all global variables of S which are subject to change by S , i.e.,

$\{x\} = \text{change}(S)$, and v is the list of all other global variables of S (read-only variables). (Hoare actually used a slightly different notation that is now obsolete.)

Legal procedure calls are of the form $\text{call } p(a : e)$, where

- a is a list of distinct variables of the same length as x that are substituted for x ,
- e is a list of expressions not containing any variable of a , of the same length as v , that are substituted for v .

The following proof rule dealt with a ‘generic’ procedure call $\text{call } p(x : v)$:

RECURSION

$$\frac{\{P\} \text{call } p(x : v) \{Q\} \vdash \{P\} S \{Q\}}{\{P\} \text{call } p(x : v) \{Q\}}$$

where the procedure p is declared by $\text{proc } p(x : v) : S$.

(Hoare actually included the procedure declaration as an additional premise of the rule.) What is the intuition behind this rule? Hoare states in [Hoa71a] that it permits

“the use of the desired conclusion as a hypothesis in the proof of the body itself.”

More specifically, the symbol \vdash in the premise denotes the provability relation. So this rule is actually a metarule. According to [FH71] the premise of this rule

“permits $\{P\} \text{call } p(x : v) \{Q\}$ to be assumed as a hypothesis in the proof of $\{P\} S \{Q\}$.”

This proof is supposed to be carried out using the remaining axioms and proof rules. The conclusion of the rule then coincides with this hypothesis.

To transfer a result established by the recursion rule to any other procedure call with actual parameters, say the lists a and e , the following substitution rule was introduced:

SUBSTITUTION

$$\frac{\{P\} \text{call } p(x : v) \{Q\}}{\{P[k := k', x := a, v := e]\} \text{call } p(a : e) \{Q[k := k', x := a, v := e]\}}$$

where the following holds for the substitutions applied to P and Q :

- k is a list of free variables of P or Q that do not occur in x or v , but which occur in a or e ,
- k' is a list of fresh variables of the same length as k that are substituted for k . So the variables in k' do not appear globally in the procedure body S , are not free in P or Q , and do not occur in a or e ,
- a and e are such that the call $\text{call } p(a : e)$ is legal.

So the substitution $[x := a, v := e]$ of the formal parameters by the actual ones is carried out together with an appropriate renaming $[k := k']$ of the ‘potentially conflicting’ variables in P and Q .

Hoare noted that the above two rules are not sufficient to reason about recursive procedures. To have a more powerful proof method, he introduced the following rule, where $\text{free}(P)$ stands for the set of free variables in an assertion P and similarly with $\text{free}(P, Q)$. Further, for a list z of variables, $\{z\}$ denotes the *set* of all variables occurring in z . Similarly, $\{a, e\}$ denotes the set of all variables occurring in the lists a and e .

ADAPTATION

$$\frac{\{P\} \text{call } p(a : e) \{Q\}}{\{\exists z (P \wedge \forall a (Q \rightarrow R))\} \text{call } p(a : e) \{R\}}$$

where z is a list of variables with $\{z\} = \text{free}(P, Q) \setminus (\text{free}(R) \cup \{a, e\})$.

The precondition of the conclusion of this rule looks complicated. What does it express? Hoare explained in [Hoa71a]:

“If R is the desired result of executing a procedure call, $\text{call } p(a : e)$, and $\{P\} \text{call } p(a : e) \{Q\}$ is already given, what is the weakest precondition W such that $\{W\} \text{call } p(a : e) \{R\}$ is universally valid? It turns out that this precondition is $\exists z (P \wedge \forall a (Q \rightarrow R))$.”

We shall discuss this rule further in Sect. 5.1.

To deal with the declarations of local variables Hoare introduced the following rule:

DECLARATION

$$\frac{\{P\} S[x := y] \{Q\}}{\{P\} \mathbf{begin\ var\ } x; S \mathbf{end\ } \{Q\}}$$

where $y \notin \mathit{free}(P, Q)$ and y does not appear in S unless the variables x and y are the same.

Additionally, a proof rule, originally proposed in [Lau71], to reason about the conditional **if-then** statement was used. We present here instead a more common rule that deals with the **if-then-else** statement:

CONDITIONAL

$$\frac{\{p \wedge B\} S_1 \{q\}, \{p \wedge \neg B\} S_2 \{q\}}{\{p\} \mathbf{if\ } B \mathbf{ then\ } S_1 \mathbf{ else\ } S_2 \mathbf{ fi\ } \{q\}}$$

The correctness proof of Quicksort by Foley and Hoare in [FH71] was carried out using the above proof rules, originally presented in [Hoa71a]. The authors formulated two correctness criteria that should hold upon termination of Quicksort:

- *Sorted*: the output array should be sorted within the given bounds m and n .
- *Perm*: the output array should be a permutation of the original input array within the given bounds m and n but untouched outside these bounds.

The proof established these properties simultaneously, using appropriate assertions. In [AdBO09] a detailed modular correctness proof of Quicksort was given. Modular means here that the property *Perm* was proved first and next, based on this result, the property *Sorted*.

3.3. Reasoning about termination

So far we did not discuss within Hoare's logic the subject of program termination. Nowadays, one talks of *partial correctness*, which refers to the conditional statement 'if the program terminates starting from a given precondition, then it satisfies the desired postcondition' and this is precisely what Hoare's proof system allows one to accomplish. A more demanding property is *total correctness*, which stipulates that all program computations starting from a given precondition terminate and satisfy the desired postcondition. To formalize these notions we need to refer to the program semantics. We shall discuss it in the next section.

All approaches to proving program termination within Hoare's logic formalize Floyd's [Flo67] observation that

"Proofs of termination are dealt with by showing that each step of a program decreases some entity which cannot decrease indefinitely."

The challenge is to incorporate such a reasoning into Hoare's framework in a simple way. The first extension of Hoare's proof system to total correctness was proposed in [MP74], but the proposed strengthening of the WHILE rule was somewhat elaborate. In [Har79] the appropriate rule took a simpler form:

WHILE I

$$\frac{P(n+1) \rightarrow B, \{P(n+1)\} S \{P(n)\}, P(0) \rightarrow \neg B}{\{\exists n P(n)\} \mathbf{while\ } B \mathbf{ do\ } S \mathbf{ od\ } \{P(0)\}}$$

where $P(n)$ is an assertion with a free variable n that does not appear in S and ranges over natural numbers.

Still, a disadvantage of this rule is that it requires to find a parameterized loop invariant $P(n)$ such that the value of n decreases exactly by 1 with each loop iteration. Such a precise information is not needed to establish termination and sometimes is difficult to come up with.

Additionally, it is often inconvenient to reason about partial correctness and termination at the same time. These concerns were addressed in the following proof rule introduced in [OG76a] that adds two new premises to the original WHILE rule:

WHILE II

$$\frac{\begin{array}{l} \{P \wedge B\} S \{P\}, \\ \{P \wedge B \wedge t = z\} S \{t < z\}, \\ P \rightarrow t \geq 0 \end{array}}{\{P\} \mathbf{while} B \mathbf{do} S \mathbf{od} \{P \wedge \neg B\}}$$

where t is an integer expression, called a *termination function* (sometimes called a *bound function* or a *variant*), and z is an integer variable that does not appear in P , B , t or S .

This proof rule corresponds to Dijkstra's modification of his weakest precondition semantics proposed in [Dij76b] and reproduced as [Dij82]. Returning to the DIV program from Sect. 3.1 note that it does not terminate when $y = 0$. To prove its termination one needs to assume that initially $x \geq 0 \wedge y > 0$ and use a stronger loop invariant, namely $P' \equiv r \geq 0 \wedge y > 0 \wedge x = r + y \cdot q$. The termination function is particularly simple here: it is just r . The relevant claims,

$$\{P' \wedge y \leq r \wedge r = z\} r := r - y; \quad q := 1 + q \{r < z\}$$

and

$$P' \rightarrow r \geq 0,$$

are straightforward to prove.

Hoare did discuss program termination in the already discussed [Hoa71b], where he showed termination of the FIND program. Since this property is not captured by his proof system from [Hoa69], he used informal arguments. According to the above terminology, he established total correctness of the program FIND. Hoare noticed that the termination proof required new invariants in addition to those needed for proving partial correctness. However, he did not introduce the concept of a termination function with a corresponding proof rule for total correctness of **while** programs.

To deal with the total correctness of the recursive procedures the following analogue of the WHILE I rule was proposed independently in [Cla76] and [Sok77]:

RECURSION I

$$\frac{\neg P(0), \{P(n)\} \mathbf{call} p \{Q\} \vdash \{P(n+1)\} S \{Q\}}{\{\exists n P(n)\} \mathbf{call} p \{Q\}}$$

given the procedure declaration **proc** $p : S$, and where $P(n)$ is an assertion with a free variable n that does not appear in S and ranges over natural numbers.

This rule shares with the WHILE I rule the same disadvantage concerning the parameterized assertion $P(n)$. This matter was addressed in [AdBO09], where the following analogue of the WHILE II rule was used for recursive procedures:

RECURSION II

$$\frac{\begin{array}{l} \{P \wedge t < z\} \mathbf{call} p \{Q\} \vdash \{P \wedge t = z\} S \{Q\} \\ P \rightarrow t \geq 0 \end{array}}{\{P\} \mathbf{call} p \{Q\}}$$

given the procedure declaration **proc** $p : S$, where t is a termination function, z is an integer variable that does not occur in P , t , Q and S and is treated in the proofs as a constant.

The last restriction means that in the proof in the first premise no proof rules are apply quantification or substitutions to the variable z . Examples of such rules are introduced in Sect. 5.1. We shall discuss the need for such a restriction in Sect. 6.1.

In the presentation of the correctness of the recursive procedure Quicksort in [FH71] only few remarks were spent on termination. In [AdBO09] termination of this program was proved using an extension of the RECURSION II rule to procedures with the call-by-value parameter mechanism.

4. Soundness and Completeness Matters

4.1. Preliminaries

In mathematical logic a standard way to judge the adequacy of a proof system is by means of the soundness and completeness concepts. It is then legitimate to address these matters for the proof systems introduced in the previous section. This requires some care since the CONSEQUENCE rule also uses implications between assertions as premises, the WHILE II rule refers to integer variables and expressions, whereas the RECURSION rules refer in their premises to the provability relation. For these considerations one needs to define some semantics with respect to which the introduced axioms and proof rules can be assessed.

To proceed in a systematic way we need to recall some basic notions from mathematical logic. Assume a first-order language \mathcal{L} . An *interpretation* I for \mathcal{L} consists of

- a non-empty domain D ,
- an assignment to each n -ary function symbol in \mathcal{L} of an n -ary function over D ,
- an assignment to each n -ary predicate symbol \mathcal{L} of an n -ary relation over D .

Given an interpretation I , a *state* σ is a function from the set of variables to the domain D . The definition of I disregards our assumption that all variables are typed. However, it is easy to amend it by replacing the domain D by the set of typed domains and by stipulating that each variable ranges over the domain associated with its type.

The next step is to define, given an interpretation I and a state σ , the value $\sigma(t)$ of an expression and when σ satisfies a formula ϕ of \mathcal{L} , written as $\sigma \models_I \phi$. We omit both definitions. We then say that a formula ϕ is *true* (or *valid*) in I , written as $\models_I \phi$, if for all states σ we have $\sigma \models_I \phi$.

The consecutive step is to define semantics of the underlying programming concepts. This can be done in a number of ways. The common denominator of all approaches is the above concept of a state. Assume that all function and predicate symbols of the considered programming languages belong to the assumed first-order language. (In fact, this assumption is already implicit in the way assertions are formed in the ASSIGNMENT axiom and the WHILE rule.) However, some adjustments are needed to deal with more complex expressions that are customary in programming languages. In particular, we need to extend the domain of the states to array variables. We stipulate that to such variables a state assigns a function from the domain of the array to the range type. Using such a function we can then assign values to expressions that use array variables. As the complexity of the considered programming language grows, its expressions become more involved and the concept of the state gets more complex. At this stage we limit ourselves to the notion of a state that assigns values to all simple variables (i.e. variables of a simple type) and array variables.

The final step is to define *semantics of the programs*. Several approaches were proposed in the literature. Their discussion and comparison is beyond the scope of this paper. For the sake of the subsequent discussion we assume a semantics of the programs that allows us to define *computations* of each considered program, which are identified here with the sequences of states that can be generated by it.

In Hoare's logic the types of the variables in the considered programs, for instance in the program DIV studied in Fig. 4, are usually omitted and one simply assumes that all variables are typed and that the considered programs are correctly typed.

4.2. Soundness

Let us return now to assertions and programs. Suppose that all assertions are formulas in a given first-order language \mathcal{L} and that all considered programs use function and predicate symbols of \mathcal{L} . Each interpretation I for \mathcal{L} then determines the set of states and thus allows us for each program to define the set of its computations over I . This in turn allows us to introduce the following notions.

Let I be an interpretation. We say that the correctness formula $\{P\} S \{Q\}$ is true in I in the sense of *partial correctness* if the following holds:

every terminating computation of S over I that starts in a state that satisfies P ends in a state that satisfies Q .

Further, we say that the correctness formula $\{P\} S \{Q\}$ is true in I in the sense of *total correctness* if the following holds:

every computation of S over I that starts in a state that satisfies P terminates and ends in a state that satisfies Q .

Denote now by \mathcal{H} the original proof system of Hoare presented in Sect.3.1 and by \mathcal{HT} the proof system obtained from \mathcal{H} by replacing the WHILE rule by the WHILE II rule. The following two results capture the crucial properties of these proof systems.

Soundness Theorem 1 Consider a proof of the correctness formula $\{P\} S \{Q\}$ in the proof system \mathcal{H} that can use in the applications of the CONSEQUENCE rule assertions from a set \mathcal{A} . Consider an interpretation I in which all assertions from \mathcal{A} are true. Then $\{P\} S \{Q\}$ is true in I in the sense of partial correctness.

This property of the proof system \mathcal{H} is called soundness in the sense of *partial correctness*. It was first established in [HL74] w.r.t. the relational semantics in which programs were represented as binary relations on the sets of states. Returning to the original proof of [Hoa69] presented in Fig. 4, the appropriate set \mathcal{A} of the assertions consists of the formulas listed in lines 1 and 6. Since they are true in the standard interpretation I of Peano arithmetic, we conclude by the above theorem that the established correctness formula listed in line 12 of Fig. 4 is true in I in the sense of partial correctness.

The following counterpart of the above theorem justifies the reasoning about termination. It is, however, important to read it in conjunction with the qualifications that follow.

Soundness Theorem 2 Consider a proof of the correctness formula $\{P\} S \{Q\}$ in the proof system \mathcal{HT} that can use in the applications of the CONSEQUENCE rule assertions from a set \mathcal{A} . Consider an interpretation I in which all assertions from \mathcal{A} are true. Then $\{P\} S \{Q\}$ is true in I in the sense of total correctness.

This property of the proof system \mathcal{HT} is called soundness in the sense of *total correctness*. The first proof was given in [Har79] and referred to the proof system in which instead of the WHILE II rule the WHILE I rule was used. In this rule the assertion $P(n)$ refers to a free variable n that ranges over natural numbers. To guarantee the correct interpretation of such assertions one needs to ensure that in each state such a variable n is interpreted as a variable of type ‘natural number’. In [Har79] this is achieved by considering assertion languages that extend the language of Peano arithmetics and by limiting one’s attention to *arithmetic interpretations*. These are interpretations that extend the standard model for arithmetic.

In the case of the WHILE II rule similar considerations are needed to ensure the correct interpretation of the integer expression t and the integer variable z . The corresponding result was given in [AO91] and reproduced in the subsequent two editions of the book. As in [AO91] all variables are assumed to be typed, t and z are correctly interpreted and the need for the arithmetic interpretations disappears.

4.3. Completeness

The completeness of the \mathcal{H} and \mathcal{HT} proof systems aims at establishing some form of converse of the Soundness Theorems. It is a subtle matter and requires a careful analysis. Let us start with the proof system \mathcal{H} . It is incomplete for an obvious reason. Consider for instance the correctness formula $\{\mathbf{true}\} x := 0 \{x \neq 1\}$. By the ASSIGNMENT axiom we get $\{0 \neq 1\} x := 0 \{x \neq 1\}$. To conclude the proof we need to establish the obvious implication $\mathbf{true} \rightarrow 0 \neq 1$ and apply the CONSEQUENCE rule. However, we have no proof rules and axioms that allow us to derive this implication.

A way out is to augment \mathcal{H} by a proof system allowing us to prove all true implications between the assertions. Unfortunately, in general such proof systems do not exist. This is a consequence of two results in mathematical logic. The first one states that the set of theorems in a proof system with recursive sets of axioms and finitary rules is recursively enumerable. The second one is Tarski’s undefinability theorem of 1936, see [Tar36]. It implies that the set of formulas of Peano arithmetic that are true in the standard model of arithmetic is not arithmetically definable, so in particular not recursively enumerable. This means that completeness of the proof system \mathcal{H} cannot be established even if we add to it a proof system concerned with the assertions.

A natural solution is to try to establish completeness *relative* to the set of true assertions, that is to use the set of true assertions as an ‘oracle’ that can be freely consulted in the correctness proof. However, even then a

problem arises because the assertion language can fail to be sufficiently expressive. Namely [Wan78] exhibited a true correctness formula that cannot be proved because the necessary intermediate assertions cannot be expressed in the considered assertion language. Simpler examples of such assertion languages were provided in [BT82a].

A solution to these complications was proposed by S.A. Cook in [Coo78]. To explain it we need to introduce some additional notions. We call a set of states Σ *definable* in an interpretation I iff for some formula ϕ we have $\Sigma = \{\sigma \mid \sigma \models_I \phi\}$.

Next, we assign to each program S its *meaning* $\mathcal{M}_I[S]$ relative to I , defined by

$$\mathcal{M}_I[S](\sigma) = \{\tau \mid \text{there exists a computation of } S \text{ over } I \text{ that starts in } \sigma \text{ and terminates in } \tau\}.$$

At this moment the set $\mathcal{M}_I[S](\sigma)$ has at most one element, which will not be anymore the case when nondeterministic or parallel programs are considered.

Then given an assertion P and a program S we define

$$sp_I(P, S) = \{\tau \mid \exists \sigma (\sigma \models_I P \wedge \tau \in \mathcal{M}_I[S](\sigma))\}.$$

So $sp_I(P, S)$ is the set of states that can be reached by executing S over I starting in a state satisfying P ; '*sp*' stands for the *strongest postcondition*.

We then say that the language \mathcal{L} is *expressive* relative to an interpretation I and a class of programs \mathcal{S} if for every assertion P and program $S \in \mathcal{S}$ the set of states $sp_I(P, S)$ is definable. Finally, given a first-order language \mathcal{L} , a proof system PS for a class of programs \mathcal{S} is called *complete in the sense of Cook* if for every interpretation I such that \mathcal{L} is expressive relative to I and \mathcal{S} the following holds:

every correctness formula true in I in the sense of partial
correctness can be proved in PS assuming all true formulas in I .

In other words, completeness in the sense of Cook is a restricted form of relative completeness mentioned above, in which we limit ourselves to the class of interpretations w.r.t. which the underlying language \mathcal{L} is expressive. The following result is a special case of a theorem established in [Coo78].

Completeness Theorem The proof system \mathcal{H} for partial correctness of the **while** programs is complete in the sense of Cook.

The main difficulty in the proof, that proceeds by induction on the program structure, consists in finding the loop invariants. A simpler argument was provided in [Cla79], where a dual definition of expressiveness was used. Instead of the strongest postcondition it relied on the so-called *weakest liberal precondition*, which, given an interpretation I , assertion Q and a program S , is defined by

$$wlp_I(S, Q) = \{\sigma \mid \forall \tau (\tau \in \mathcal{M}_I[S](\sigma) \rightarrow \tau \models_I Q)\}.$$

So $wlp_I(S, Q)$ is the set of states from which all terminating computations of S over I end in a state satisfying Q . The qualification 'liberal' refers to the fact that termination is not guaranteed. The assumption that the set of states $wlp_I(Q, S)$ is definable makes it possible to find a very simple loop invariant. Namely, assuming that $\{P\}$ **while** B **do** S **od** $\{Q\}$ is true in an interpretation I such that \mathcal{L} is expressive relative to it in this revised sense, it turns out that a loop invariant is simply an assertion R defining $wlp_I(\mathbf{while } B \mathbf{ do } S \mathbf{ od}, Q)$. Additionally, both $P \rightarrow R$ and $R \wedge \neg B \rightarrow Q$ are true in I , which allows one to establish $\{P\}$ **while** B **do** S **od** $\{Q\}$ by the WHILE and CONSEQUENCE rules.

The above completeness result led to a number of works that tried to clarify it in a logical setting. In particular, J.A. Bergstra and J.V. Tucker noticed in [BT82b] that expressiveness is not a necessary condition for a Hoare's logic to be complete. Indeed, if one takes as the assertion language the language of Peano arithmetic, then there exist interpretations I for which it is not expressive, while all correctness formulas that are true in I in the sense of partial correctness can be proved in the proof system \mathcal{H} .

We should also mention here a proposal put forward in [BG87] by A. Blass and Y. Gurevich. They suggested to use a different assertion language than first-order logic (or its multi-sorted variants dealing with subscripted variables or typed variables). The proposed assertion language is a fragment of the second-order logic, called *existential fixed-point logic* (EFL). EFL extends a fragment of first-order logic, in which negation is applied only to atomic formulas and the universal quantifier is absent, by a fixed-point operator. The authors showed that EFL is sufficient for proving relative completeness of the proof system \mathcal{H} without any expressiveness assumption.

The reason is that both the strongest postconditions and the weakest liberal preconditions of the **while** programs (also in presence of recursive parameterless procedures) are definable in EFL.

Different insights on the relative completeness can be obtained by just limiting the assertion language to propositions. The resulting propositional fragment of Hoare's logic of **while** programs (PHL) can then be studied in a purely algebraic way, by means of an equational reasoning. D. Kozen showed in [Koz00] that an extension of Kleene's algebra of regular expressions with tests (KAT) subsumes PHL. In this setting the **while** B **do** S **od** program is represented by the expression $(BS)^* \bar{B}$ and the correctness formula $\{P\} S \{Q\}$ by the equation $PS\bar{Q} = 0$, where $\bar{}$ is the complementation operator of KAT. After this translation the proof rules of the proof system \mathcal{H} become theorems of KAT. Further, after defining validity of implications between conjunctions of expressions using a relational semantics, a counterpart of the completeness result can be established for PHL. This shows that reasoning within propositional Hoare's logic can be represented by purely equational reasoning. This work was followed by [KT01], where it was clarified how the notion of completeness in the sense of Cook can be captured by adding the weakest liberal precondition to the assertion language and an appropriate axiomatization of it to the axioms of KAT.

An even more general analysis of soundness and relative completeness of the proof system \mathcal{H} was presented in [AMMO09] within the framework of category theory, using a simple theory of pre-ordered sets. Such a general approach allowed the authors to derive new sound and relatively complete Hoare's logics for the runtime analysis and for verification of linear dynamical systems.

Consider now the proof system \mathcal{HT} for total correctness of **while** programs. To establish its completeness in the appropriate sense we encounter the same complications as in the case of \mathcal{H} , but additionally we have to deal with the problem of definability of the termination functions used in the WHILE II rule. In [Har79] completeness was established for the assertion languages that extend the language of Peano arithmetic and for arithmetic interpretations defined in the previous subsection, but the paper considered the WHILE I rule in which the termination functions are absent. To be able to express in the assertion language the parameterized loop invariants $P(n)$ the definition of an arithmetic interpretation stipulates that there is a formula in the assertion language that, when interpreted, encodes finite sequences of the domain elements by one element.

In [AO91] and the subsequent two editions of the book relative completeness of \mathcal{HT} was established. To this end, it was assumed that the underlying assertion language is *expressive*, which meant that for every **while** loop S there exists an integer expression t such that whenever S terminates when started in a state σ , then the value $\sigma(t)$ is the number of loop iterations. In the adopted setup the assumption that all variables are typed automatically ensures that the considered interpretations included the standard model of Peano arithmetic and that $\sigma(t)$ is a natural number.

5. Fine-tuning the Approach

The matters discussed until now gloss over certain issues that have to do with the adjustments of the preconditions and postconditions, various uses of variables, and procedure parameters. In this section we discuss closely these matters, as they reveal some differences between customary logics and Hoare's logic and show the subtleties of reasoning about various uses of variables in the context of programs.

5.1. Adaptation rules

In Hoare's logic we see two types of rules. First, for each programming construct there is at least one axiom or rule dealing with its correctness. Together, they make possible a syntax-directed reasoning about program correctness. Second, there are proof rules where the same program S is considered in the premise and the conclusion. These rules allow us to *adapt* an already established correctness formula $\{P_1\} S \{Q_1\}$ about S to another proof context. Most prominent is the CONSEQUENCE rule that allows us to strengthen the precondition P_1 to a precondition P with $P \rightarrow P_1$ and to weaken the postcondition Q_1 to a postcondition Q with $Q_1 \rightarrow Q$, thus arriving at the conclusion $\{P\} S \{Q\}$. Another one is Hoare's ADAPTATION rule dealing with procedure calls. Hoare stated in [Hoa71a] that in the absence of recursion, i.e., in his proof system for **while** programs, his ADAPTATION rule is a derived rule. So the power of this rule is only noticeable in the context of recursion.

Other rules can be conceived that are concerned with the same program in the premise and conclusion. For example, the following rules were introduced in [Har79] and [AdBO09]. Here and elsewhere we denote the set of variables of a program S by $var(S)$.

INVARIANCE

$$\frac{\{R\} S \{Q\}}{\{P \wedge R\} S \{P \wedge Q\}}$$

where $free(P) \cap change(S) = \emptyset$.

 \exists -INTRODUCTION

$$\frac{\{P\} S \{Q\}}{\{\exists x : P\} S \{Q\}}$$

where $x \notin var(S) \cup free(Q)$.

SUBSTITUTION I

$$\frac{\{P\} S \{Q\}}{\{P[\mathbf{z} := \mathbf{t}]\} S \{Q[\mathbf{z} := \mathbf{t}]\}}$$

where $(\{\mathbf{z}\} \cap var(S)) \cup (var(\mathbf{t}) \cap change(S)) = \emptyset$.

We shall return to these rules shortly. But first, following [Old83b], let us discuss the ADAPTATION rule in the more general setting of programs. We say that a program S is *based on* a finite set X of variables if $var(S) \subseteq X$ holds. Now we can recast Hoare's ADAPTATION rule as follows:

ADAPTATION I

$$\frac{\{P\} S \{Q\}}{\{\exists \mathbf{z}(P \wedge \forall \mathbf{x}(Q \rightarrow R))\} S \{R\}}$$

where \mathbf{x} and \mathbf{z} are lists of variables, S is based on $\{\mathbf{x}\}$, and $\{\mathbf{z}\} = free(P, Q) \setminus (free(R) \cup \{\mathbf{x}\})$.

Following Hoare, the precondition in the conclusion of this rule intends to express the weakest precondition W such that $\{W\} S \{R\}$ holds (in the sense of partial correctness), assuming that R is the desired result of executing S and $\{P\} S \{Q\}$ is already established. This intention can be phrased as follows: find the weakest assertion W such that $\{W\} S \{R\}$ holds for *all* programs based on $\{\mathbf{x}\}$ that satisfy $\{P\} S \{Q\}$. In [Old83b] this precondition was calculated as follows:

$$W \equiv \forall \mathbf{y} (\forall \mathbf{u} (P \rightarrow Q[\mathbf{x} := \mathbf{y}]) \rightarrow R[\mathbf{x} := \mathbf{y}]),$$

where \mathbf{y} is a list of fresh variables of the same length as \mathbf{x} with $\{\mathbf{y}\} \cap free(P, Q, R) = \emptyset$ and $\{\mathbf{u}\} = free(P, Q) \setminus \{\mathbf{x}\}$.

Comparing W with the precondition used in the conclusion of the ADAPTATION I rule shows that the implication

$$\exists \mathbf{z}(P \wedge \forall \mathbf{x}(Q \rightarrow R)) \rightarrow W$$

holds but the converse is false. Thus Hoare's precondition is sound but is stronger than necessary. This suggests the following variant of the rule:

ADAPTATION II

$$\frac{\{P\} S \{Q\}}{\{W\} S \{R\}}$$

where S is based on $\{\mathbf{x}\}$ and W is the precondition calculated above.

To show the weakness of the Hoare's ADAPTATION I rule in the context of recursive procedures, consider the following example essentially taken from [Old83b]. Let I be the standard interpretation of integers and consider the procedure declaration

```
proc p(x :) : if x > 1 then x := x - 1; call p(x :); x := x + 1 fi,
```

where x is an integer variable. Then the procedure call satisfies the following correctness formula

$$\{x = y\} \text{ call } p(x :) \{x = y\},$$

where y is a fresh variable x . It expresses the fact that the procedure call leaves the variable x unchanged. For the postcondition $x = y + 1$ the ADAPTATION I rule (or, equivalently, the original ADAPTATION rule) permits us to conclude only

$$\{\text{false}\} \text{ call } p(x :) \{x = y + 1\},$$

whereas the weakest liberal precondition $wlp_I(\text{call } p(x :), x = y + 1)$ is expressed by the formula $x = y + 1$. This precondition is correctly calculated by the ADAPTATION II rule, which yields

$$\{x = y + 1\} \text{ call } p(x :) \{x = y + 1\}.$$

To compare the power of different adaptation rules in a systematic way, S. de Gouw and J. Rot [dGR16] used the following notion due to [Kle99]. A set \mathcal{R} of proof rules for a class \mathcal{S} of programs is called *adaptation complete* if for all assertions P, Q, P', Q' and finite sets X of variables

- whenever for all programs $S \in \mathcal{S}$ based on X the truth of $\{P\} S \{Q\}$ implies the truth of $\{P'\} S \{Q'\}$ in the sense of partial correctness,
- then for all program $S \in \mathcal{S}$ based on X there is a derivation of $\{P'\} S \{Q'\}$ from $\{P\} S \{Q\}$ using only rules of \mathcal{R} , written as $\{P\} S \{Q\} \vdash_{\mathcal{R}} \{P'\} S \{Q'\}$.

By the result of [Old83b], the set $\mathcal{R}_O = \{\text{ADAPTATION II, CONSEQUENCE}\}$ is adaptation complete. Further, \mathcal{R}_O enjoys two properties, as noted in [dGR16]:

1. Other adaptation rules, like \exists -INTRODUCTION and slightly weakened versions of the INVARIANCE and SUBSTITUTION I rules are derivable from \mathcal{R}_O .
2. Any derivation in \mathcal{R}_O can be replaced by a single application of each of the two rules in \mathcal{R}_O .

What about Hoare's adaptation rule? Let $\mathcal{R}_H = \{\text{ADAPTATION I, CONSEQUENCE}\}$. From the above example it follows that this set is *not* adaptation complete. Nevertheless, \mathcal{R}_H enjoys property 1, but not property 2 of \mathcal{R}_O .

The paper [Old83b] also investigated three other adaptation rules proposed in the literature. An adaptation rule introduced in [GL80] turned out to be sound but *not* adaptation complete when grouped together with the CONSEQUENCE rule. In turn an adaptation rule for the programming language EUCLID given in [LGH⁺78] is not even sound, while an adaptation rule introduced in [CO81] is both sound and adaptation complete when grouped together with the CONSEQUENCE rule.

5.2. Subscripted and local variables

Subscripted variables In both [Hoa71b] and [FH71] programs with arrays were studied. To reason about assignments to the subscripted variables the ASSIGNMENT axiom was used, implicitly assuming that the definition of substitution to subscripted variables is obvious. This is indeed the case when the array subscripts are simple expressions, for example a constant or a simple variable, which was indeed the case for both programs analyzed there. However, in the case of more complex subscripts difficulties may arise, as the following example discussed in [dB80] shows. In the case of an assignment to a simple variable any correctness formula $\{P\} x := t \{x = t\}$, where t is a constant is true. However, the correctness formula

$$\{a[1] = 2 \wedge a[2] = 2\} a[a[2]] := 1 \{a[a[2]] = 1\}$$

is false. Indeed, given the precondition the execution of the assignment $a[a[2]] := 1$ amounts to executing the assignment $a[2] := 1$ after which the expression $a[a[2]]$ evaluates to 2 and not 1. This suggests that the ASSIGNMENT axiom cannot be used for arbitrary subscripted variables.

This complication was clarified and solved in a systematic way in [dB80], by extending the definition of substitution to an arbitrary subscripted variable. The crucial step in the inductive definition of the substitution $s[u := t]$ deals with the case when $s \equiv a[s_1]$ and $u \equiv a[u_1]$, for which one defines

$$s[u := t] \equiv \text{if } s_1[u := t] = u_1 \text{ then } t \text{ else } a[s_1[u := t]] \text{ fi.}$$

So in the **if** case one checks whether after performing the substitution $[u := t]$ on s_1 the subscripts s_1 and u_1 are *aliases*—and substitutes in that case $a[s_1]$ by t — while in the **else** case one applies the substitution $[u := t]$ inductively to the subscript s_1 of $a[s_1]$.

J.W. de Bakker showed that with this extended definition of substitution the ASSIGNMENT axiom remains sound for subscripted variables. Different axioms for assignment to subscripted variables were given in [HW73, Gri78, Apt81]. In the first two references in the corresponding axiom an array is treated as a partial function and an assignment to a subscripted variable is interpreted as a change in the whole function.

Local variables Consider now the case of local variables. They can be viewed as a counterpart of bound variables in logical formulas. However, the situation is more complicated because of the dynamic character of variables in programming languages and the presence of procedures.

We discussed already completeness in the sense of Cook of the proof system \mathcal{H} given in [Coo78]. Cook actually considered an extension of the proof system \mathcal{H} by axioms and proof rules for a small programming language that allows variable declarations and nonrecursive procedures and proved its completeness in the above sense. However, the semantics of the block statement made the corresponding completeness result invalid. It is useful to discuss this matter more closely.

Local variables were already dealt with in the DECLARATION rule mentioned in Sect.3.2. In [Coo78] a different rule was used in which the substitution was applied to the assertions and not to the programs, namely:

BLOCK

$$\frac{\{P[x := y]\} S \{Q[x := y]\}}{\{P\} \mathbf{begin\ var\ } x; S \mathbf{end\ } \{Q\}}$$

where $y \notin \mathit{free}(P, Q) \cup \mathit{var}(S)$.

More precisely, a slightly adjusted version of this rule was used so that one could reason about variable declarations in the context of non-recursive procedures. But even without this adjustment a possible problem arises. Consider the program

begin var x ; $x := 1$ **end**; **begin var** y ; $z := y$ **end**.

In many programming languages it would yield an error because the right-hand side of the second assignment refers to a value of the uninstantiated variable y . However, according to the semantics proposed in [Coo78] such assignments were allowed. Local variables were modelled using a stack in which the last used value was kept on the stack and implicitly assigned to the next local variable. As a result the correctness formula

{true} **begin var** x ; $x := 1$ **end**; **begin var** x ; $y := x$ **end** **{ $y = 1$ }**

was true according to the semantics though there is no way to prove it.

[Coo81] provided a corrigendum in which two possible fixes were suggested. One was to modify the semantics so that the proposed proof system is still complete. This can be achieved by assigning to each newly declared variable a register that has not been used before and modifying the notion of a state accordingly.

Another one was to require all newly declared variables to be initialized to some fixed value, say ω . This option was first used in [Gor75], where a variant of the BLOCK rule was used in which the initialization was modelled on the proof-theoretic level by the conjunct $x = \omega$ added to the precondition of the premise. To correct the proof of the relative completeness result given in [Coo78], one should then use an initialization of the local variables, both in the semantics and in the proof theory, so use the variant of the BLOCK rule given in [Gor75].

Yet another option is to require all newly declared variables to be explicitly initialized to some arbitrary expression. This approach was taken in [AdBO09], where the following more general version of the corresponding rule was used that allowed a declaration of a list of new variables:

BLOCK I

$$\frac{\{P\} \mathbf{x := t}; S \{Q\}}{\{P\} \mathbf{begin\ var\ } \mathbf{x := t}; S \mathbf{end\ } \{Q\}}$$

where $\{\mathbf{x}\} \cap \mathit{free}(Q) = \emptyset$.

Here $\mathbf{x} := \mathbf{t}$, where \mathbf{x} is a list of different variables and \mathbf{t} a corresponding list of expressions, is a *parallel assignment*, introduced in [Dij75] and further discussed in Sect. 10.

It is natural to postulate in this rule that the variables listed in \mathbf{x} do not appear in the expressions from \mathbf{t} . However, this is a syntactic condition concerning the program formation that is not needed to reason about partial correctness. Further, as we shall soon see, putting no restrictions on \mathbf{x} and \mathbf{t} turns out to be useful for modelling parameter passing in a subtle situation when some formal parameters happen to coincide with the global variables that are used in actual parameters.

The observant reader will notice that in the discussed rules substitution is used differently. In the DECLARATION rule the substitution is applied to the programs, in the BLOCK rule it is applied to the assertions, while—interestingly—in the BLOCK I rule it is not used at all. The resulting proof systems yield different results when applied to programs that use procedures. To illustrate the problem consider the parameterless procedure declared by $\mathbf{proc} \ p : z := x$, the program

$$S_0 \equiv \mathbf{begin\ var} \ x; \ x := 0; \ \mathbf{call} \ p \ \mathbf{end},$$

and the correctness formula

$$\{x = 1\} S_0 \{z = 1\}. \quad (1)$$

To reason about the procedure call $\mathbf{call} \ p$ we add to the proof system \mathcal{H} the following degenerated version of the RECURSION rule:

COPY

$$\frac{\{P\} S \{Q\}}{\{P\} \mathbf{call} \ p \{Q\}}$$

assuming the declaration of a parameterless non-recursive procedure $\mathbf{proc} \ p : S$.

In our case it allows us to derive $\{x = 1\} \mathbf{call} \ p \{z = 1\}$ from $\{x = 1\} z := x \{z = 1\}$. This in turn allows us to derive

$$\{x = 1\} y := 0; \ \mathbf{call} \ p \{z = 1\}.$$

Now, applying the DECLARATION rule we get (1).

However, using the BLOCK rule we get a different conclusion. Namely, we first establish

$$\{y = 1\} x := 0; \ \mathbf{call} \ p \{z = 0\},$$

from which

$$\{x = 1\} S_0 \{z = 0\}. \quad (2)$$

follows.

Finally, if we use the BLOCK I rule, and therefore consider a slightly modified program

$$S' \equiv \mathbf{begin\ var} \ x := 0; \ \mathbf{call} \ p \ \mathbf{end},$$

we get

$$\{x = 1\} S' \{z = 0\}. \quad (3)$$

These differences have to do with the way local variables are interpreted in the presence of procedures. According to the *static* scope policy the procedures should be evaluated in the environment in which they were declared, while according to the *dynamic* scope policy they should be evaluated in the environment in which they were called. So according to the static scope, which is adopted in most imperative languages, we should conclude (1) and not (2) or (3).

We conclude that the proof systems studied in [Coo78] and [Gor75] dealt with dynamic scope and not static scope. The same is in principle the case for [AdBO09], but one requires there that the local variables are first renamed so that they differ from global variables. The effect is that then static and dynamic scopes coincide. In the above example one thus considers the statement

$$S_1 \equiv \mathbf{begin\ var} \ y; \ y := 0; \ \mathbf{call} \ p \ \mathbf{end}$$

instead of S_0 . Then we get $\{x = 1\} S_1 \{z = 1\}$, as desired.

5.3. Parameter mechanisms and procedure calls

The *call-by-name* parameter mechanism was originally proposed in ALGOL 60. It was used in [Hoa71a] and [Coo78] and adopted in all subsequently discussed papers on procedures, unless stated otherwise. It boils down to a simultaneous substitution of the actual parameters for the formal ones, so naturally it was modelled in the SUBSTITUTION rule by a straightforward substitution.

However, a most commonly used parameter mechanism is *call-by-value*. According to its semantics the actual parameters are evaluated first and subsequently their values assigned to the formal parameters. Some other parameter mechanisms were occasionally used. For example, the programming language PASCAL (see [JW75]) also allows the *call-by-variable* mechanism (also called *call-by-reference*), which is a mixture of call-by-name and call-by-value. The actual parameter has to be a variable. In case it is a subscripted variable, its index is evaluated first and the resulting subscripted variable is substituted for the formal parameter.

In [AdB77] it was proposed to model these two parameter mechanisms of PASCAL by means of an appropriate syntactic transformation of the procedure body, that was called 'syntactic application'. In what follows we use in the procedure declaration the qualification **val** to indicate call-by-value and **var** to indicate call-by-variable. Given a procedure declaration $\text{proc } p(\text{val } x, \text{var } y) : S$, so with x called by value and y called by variable, the call $p(t, v)$, where t is an expression and v a, possibly subscripted, variable, was modelled by the syntactic application $S[t, v]$ defined by

$$\begin{aligned} S[t, z] &\equiv \text{begin var } u; u := t; S[x := u, y := z] \text{end}, \\ S[t, a[s]] &\equiv \text{begin var } u_1, u_2; u_1 := t; u_2 := s; S[x := u_1, y := a[u_2]] \text{end}, \end{aligned}$$

where z is a simple variable and u, u_1, u_2 do not appear in s, t, z or S .

This leads to the following generalization of the COPY rule from the previous subsection:

CALL-BY-VALUE/CALL-BY-VARIABLE

$$\frac{\{P\} S[t, v] \{Q\}}{\{P\} \text{call } p(t, v) \{Q\}}$$

where the non-recursive procedure p is declared by $\text{proc } p(\text{val } x, \text{var } y) : S$.

In [AdBO09] this approach to call-by-value was simplified by noting that no variable renaming is needed to model it. The resulting rule, that needs to be used in conjunction with the BLOCK I rule, became:

CALL-BY-VALUE

$$\frac{\{P\} \text{begin var } x := t; S \text{end } \{Q\}}{\{P\} \text{call } p(t) \{Q\}}$$

where the non-recursive procedure p is declared by $\text{proc } p(\text{val } x) : S$.

To see how this rule correctly handles a subtle situation when a formal parameter coincides with a global variable used in an actual parameter, consider a procedure declared by

proc $p(\text{val } x) : x := x + 2; y := y + x$.

Using the BLOCK I rule we can then establish the correctness formula

$$\{y = 0 \wedge x = 1\} \text{begin var } x; x := x + 1; x := x + 2; y := y + x \text{end } \{y = 4\},$$

from which

$$\{y = 0 \wedge x = 1\} p(x + 1) \{y = 4\}$$

follows by the CALL-BY-VALUE rule. This agrees with the semantics of the call-by-value parameter mechanism. (The stronger postcondition, $y = 4 \wedge x = 1$ can be established using the axioms and proof rules introduced in the next section.) So the assignment $x := x + 1$ refers on the left-hand side to the formal parameter x and on the right-hand side to the actual parameter $x + 1$ that contains the global variable x .

In [Apt81] it was suggested that other parameter mechanisms can also be modelled by means of an appropriate syntactic application, and subsequently reasoned about within Hoare's logic. An example is the *call-by-result*

parameter mechanism of ALGOL W, a precursor of PASCAL (see [WH66]). According to it the actual parameter is either a simple or a subscripted variable. Upon termination of the call the value of the formal parameter is assigned to the actual parameter. In the case the actual parameter is a subscripted variable, its index is evaluated first. This parameter mechanism is used in conjunction with the call-by-value.

We shall return to this subject in the next section when discussing completeness of the counterparts of such rules for recursive procedures.

6. Reasoning about Arbitrary Procedures

6.1. Completeness results for recursive procedures

The relative completeness result established in [Coo78] dealt with the language considered in [Hoa71a] and Sect. 3.2, except that recursion was disallowed. To ensure soundness in the sense of partial correctness Cook stipulated that for the procedure calls $\text{call } p(\mathbf{a} : \mathbf{e})$ no variable in $(\mathbf{a} : \mathbf{e})$ different from formal parameters occurs globally in the procedure body.

Cook's work was extended to the language in which recursive procedures were allowed in the Master Thesis of G.A. Gorelick, written under the supervision of Cook. The details are only available as a technical report [Gor75]. We present the essentials for the case of a single recursive procedure $\text{proc } p(\mathbf{x} : \mathbf{v}) : S$, in line with the presentation in Sect. 3.2.

The conceptual contribution of Gorelick is the introduction of *most general formulas*. He wrote:

“The completeness result for recursive programs is then obtained by exhibiting, for each recursive procedure p , a “most general formula” α_p such that $\vdash \alpha_p$, and $\alpha_p \vdash \beta$ for all true formulas β about p .”

Given a procedure declaration $\text{proc } p(\mathbf{x} : \mathbf{v}) : S$, a *most general formula* for the procedure p is a correctness formula

$$\{\mathbf{c} = \mathbf{z}\} \text{ call } p(\mathbf{x} : \mathbf{v}) \{G\},$$

where \mathbf{c} is the list of variables that appear in the formal parameters \mathbf{x} and \mathbf{v} or have a global occurrence in S , and \mathbf{z} is a list a fresh variables (not occurring in \mathbf{x} , \mathbf{v} , or S), of the same length as \mathbf{c} that serves to *freeze* the initial values of the variables in \mathbf{c} before they are changed by S . The formula G is taken to express the strongest postcondition $sp_I(\mathbf{c} = \mathbf{z}, S)$ introduced in Sect. 4.3. Since the variables in \mathbf{c} and \mathbf{z} may appear in G as free variables, G describes the relationship between the initial values of \mathbf{c} represented by \mathbf{z} and the final values of \mathbf{c} computed by the procedure body S and represented by \mathbf{c} itself.

The crucial properties of most general formulas are captured by the following lemmas due to [Gor75].

Lemma G1 If $\{P\} \text{ call } p(\mathbf{x} : \mathbf{v}) \{Q\}$ is true in I in the sense of partial correctness then it can be derived from $\{\mathbf{c} = \mathbf{z}\} \text{ call } p(\mathbf{x} : \mathbf{v}) \{G\}$ and the set of all true formulas in I using “suitable adaptation rules”.

Lemma G2 For each procedure call $\text{call } p(\mathbf{c} : \mathbf{v})$ the most general formula $\{\mathbf{c} = \mathbf{z}\} \text{ call } p(\mathbf{x} : \mathbf{v}) \{G\}$ can be derived from the set of all true formulas in I using Lemma G1 and the RECURSION rule.

The proof of Lemma G1 is based on the following axiom and adaptation rules, proposed in [Gor75] for the case where S is a procedure call $\text{call } p(\mathbf{a} : \mathbf{e})$:

INVARIANCE

$$\{P\} S \{P\}$$

where $\text{free}(P) \cap \text{var}(S) = \emptyset$.

CONJUNCTION

$$\frac{\{P\} S \{Q\}, \{P\} S \{R\}}{\{P\} S \{Q \wedge R\}}$$

VARIABLE SUBSTITUTION

$$\frac{\{P\} S \{Q\}}{\{P[\mathbf{z} := \mathbf{t}]\} S \{Q[\mathbf{z} := \mathbf{t}]\}}$$

where

- $\{\mathbf{z}\} \cap \text{var}(S) = \emptyset$,
- if for any component t_i of the list \mathbf{t} , $\text{var}(t_i) \cap \text{var}(S) \neq \emptyset$, then the corresponding component z_i of the list \mathbf{z} satisfies $z_i \notin \text{free}(Q)$.

Using Lemmas G1 and G2, Gorelick established the following result. (The result was actually established for a system of recursive procedures, so the RECURSION rule was appropriately generalized.)

Completeness Theorem For programs with recursive procedures as defined in Sect. 3.2, the proof system \mathcal{H} extended by the RECURSION and SUBSTITUTION rules of Sect. 3.2 and the above axiom and adaptation rules is complete in the sense of Cook.

Both [Coo78] and [Gor75] followed [Hoa71a] and used the call-by-name parameter mechanism. The restrictions imposed on the actual parameters of 'legal' procedure calls were partly lifted in [CO81]. Analogous research was carried out in [dB80] for the case of the call-by-value and call-by-variable parameter mechanisms in the presence of recursive procedures. To this end, the CALL-BY-VALUE/CALL-BY-VARIABLE rule from the previous section was modified to the following rule:

RECURSION III

$$\frac{\{P_i\} \text{ call } p(t_i, v_i) \{Q_i\}, i \in \{1, \dots, n\} \vdash \{P_i\} S[t_i, v_i] \{Q_i\}, i \in \{1, \dots, n\}}{\{P_1\} \text{ call } p(t_1, v_1) \{Q_1\}}$$

where the procedure p is declared by $\text{proc } p(\text{val } x, \text{var } y) : S$ and $p(t_i, v_i)$, $i \in \{1, \dots, n\}$, are the procedure calls that appear in $p(t_1, v_1)$ and $S[t_1, v_1]$.

In [dB80] soundness and relative completeness for a corresponding proof system was proved. However, the proof was established only for the special case of a single recursive procedure, given the combinatorial explosion of the cases concerned with the relation between the actual and formal parameters. The main ideas of this proof were discussed in [Apt81].

Proof systems for recursive procedures with the call-by-value parameter mechanism were studied in a couple of other publications. In [vO99] a sound and relatively complete proof system was introduced for a programming language with local and global variables and mutually recursive procedures with the call-by-value mechanism, that allows for both static and dynamic scope. The proofs were not provided but it was mentioned that they were certified in the Isabelle theorem prover. In this work each assertion was identified with its meaning represented as a function from states to Boolean values. This approach simplifies the soundness and relative completeness proofs by avoiding various complications concerned with possible variable clashes (in particular, it obviates the use of the SUBSTITUTION rules), but makes it cumbersome to verify specific programs and deviates from the syntactic approach pursued in the literature on Hoare's logic.

Further, in [AdBOdG12] soundness and relative completeness of the proof system introduced in [AdBO09] was established. The recursive procedures with the call-by-value parameter mechanism were dealt with by an appropriate modification of the CALL-BY-VALUE rule from the previous section.

However, in all these works on the call-by-value and call-by-variable parameter mechanisms correctness of each procedure call had to be dealt with separately. This is an obvious drawback, since the resulting correctness proofs are then not linear in the length of the program. It would be preferable if could establish a desired property for a 'generic call' just once, from which the needed properties of all specific calls would follow by some form of a substitution rule. In [AdBO09] it was observed that this can be achieved for the calls in which no actual parameter contains a global variable. Recently, K.R. Apt and F.S. de Boer introduced in [AdB19] a sound and relatively complete proof system that provides a satisfactory solution to this problem, without any restriction. They used a generalization of the following recursion rule to a system of mutually recursive procedures:

RECURSION IV

$$\frac{\{P\} p(\mathbf{x}) \{Q\} \vdash \{P\} S \{Q\}}{\{P\} p(\mathbf{x}) \{Q\}}$$

where the procedure p is declared by $\text{proc } p(\text{val } \mathbf{x}) : S$ and $\{\mathbf{x}\} \cap \text{free}(Q) = \emptyset$.

It was combined with the following rule:

PROCEDURE CALL

$$\frac{\{P\} p(\mathbf{x}) \{Q\}}{\{P[\mathbf{x} := \mathbf{t}]\} p(\mathbf{t}) \{Q\}}$$

where the side conditions are the same as in the previous rule.

The resulting proof system uses the BLOCK I rule from the previous section, so it is appropriate for the dynamic scope. However, the authors pointed out that for a natural class of programs in which the global variables that occur in procedure bodies differ from the local variables, static and dynamic scopes coincide. Such a class of programs was first considered in [CO81].

So far we discussed partial correctness of programs with recursive procedures. In [Apt81] it was stated without a proof that a proof system corresponding to the one used in [Gor75] is sound in the sense of total correctness. In this system the INVARIANCE axiom was dropped (it is not sound for total correctness), the procedures had no parameters and the RECURSION rule was replaced by the RECURSION I rule. However, it was discovered in [AdB90a] that this claim is false. The problem has to do with the fact that the counter variable n can be subject to quantifier elimination in the \exists -INTRODUCTION rule, and to substitution in the SUBSTITUTION rule.

For example, given the procedure declaration $\text{proc } p : p; p$ of an obviously nonterminating procedure one can establish the premises $\neg P(0)$ and $\{P(n)\} \text{ call } p \{\text{true}\} \vdash \{P(n+1)\} S \{\text{true}\}$ of the above rule for $P(n) \equiv n > 1$ and conclude $\{\text{true}\} \text{ call } p \{\text{true}\}$ by the RECURSION I and CONSEQUENCES rules.

The solution proposed in [AdB90a] was to stipulate that the counter variables are treated as constants in the \exists -INTRODUCTION and SUBSTITUTION rules. This allowed the authors to prove both soundness and relative completeness of the resulting proof system for total correctness of recursive procedures without parameters w.r.t. the arithmetic interpretations introduced in Sect. 4.2. The above remarks explain why we added in the RECURSION II rule a qualification concerning the counter variable z .

6.2. Clarke's incompleteness result

Programming languages like ALGOL 60 [NBB⁺63] and PASCAL [JW75] contain procedure mechanisms that are considerably more complex than what we discussed so far. In his seminal paper [Cla79], E.M. Clarke identified a combination of features for which it is impossible to obtain a Hoare-like proof system that is sound and relatively complete (in the sense of Cook). By a Hoare-like proof system we mean a set of syntax-directed proof rules for correctness formulas about programs in the programming language under consideration such that the application of each proof rule is decidable. Moreover, the proof system allows one to prove correctness formulas concerning programs from a given programming language.

Clarke proved this incompleteness result for a block-structured programming language L_0 which includes the following features that are present in ALGOL 60 and PASCAL:

1. procedure names as parameters of procedure calls,
2. recursion,
3. static scope,
4. global variables in procedure bodies,
5. nested procedure declarations.

Clarke's incompleteness result is based on the following two lemmas for Hoare-like proof systems H and programming languages L .

Lemma A If H is sound and relatively complete for L , then the divergence problem for L is decidable for finite interpretations I .

Lemma B If L has a rich procedure concept including features 1–5, then its divergence problem is undecidable for all finite interpretations I with at least two domain elements.

An interpretation I is called *finite* if its domain D is finite. A program S in L *diverges* for I if S never terminates when started with any input values from D for its variables. The *divergence problem* for a programming language L for an interpretation I is the problem of deciding whether an arbitrary program S of L *diverges*.

The proof of Lemma A rests on the following general observations. A program S of L diverges for I if and only if the correctness formula $\{true\} S \{false\}$ is true in I in the sense of partial correctness. Since H is sound and relatively complete, the latter is true if and only if $\{true\} S \{false\}$ is provable in H , when all assertions in the proofs are interpreted in I . Thus all diverging programs S in L can be recursively enumerated by enumerating all proofs in H , thereby deciding in I the assertions used in the applications of the CONSEQUENCE rule, which is possible because I is finite. Trivially, also all non-diverging programs S in L can be recursively enumerated: simply start each program S for all of its finitely many inputs from I and enumerate it in case it halts for one of these inputs. Decidability of the set of all diverging programs follows since both the set and its complement are recursively enumerable.

For the proof of Lemma B, Clarke shows that *queue machines* (or *Post machines*), which have an undecidable halting problem, see, e.g., [Man74], can be simulated by programs from L_0 . A queue machine manipulates a queue and finitely many registers. A machine program is a finite sequence of labelled instructions of three types:

- enqueue x that adds the value of register x to the rear of the queue,
- dequeue x that removes the front entry from the queue and puts into register x , and
- if $x=y$ then go to ℓ that branches to the instruction labelled with ℓ if the values of the registers x and y agree.

Clarke [Cla79] described his simulation idea of a queue machine with a program in L as follows:

“The queue is represented by successive activations of a recursive procedure *sim* with the queue entries being maintained as values of the variable *top* which is local to *sim*. Thus an addition to the rear of the queue may be accomplished by having *sim* call itself recursively. Deletions from the front of the queue are more complicated. *sim* also contains a local procedure *up* which is passed as a parameter during the recursive call which takes place when an entry is added to the rear of the queue. In deleting an entry from the front of the queue, this parameter is used to return control to previous activations of *sim* and inspect the values of *top* local to those activations. The first entry in the queue will be indicated by marking (e.g. negating) the appropriate copy of *top*.”

From this description it is clear that the simulation program exploits the features 1, 2 and 5. Feature 4 is also needed since a global variable *program_counter* is used in the body of the procedure *sim*. Feature 3 (static scope) concerns the semantics of procedures and is needed to achieve the correct back pointers for the procedure *up* in the runtime stack generated by the successive activations of the procedure *sim*.

Note that the procedure *sim* has a formal parameter that is instantiated with procedures, which in turn have no formal procedure parameters on their own. Here the concept of mode depth is helpful. The *mode depth* of a procedure is defined inductively. A procedure that may not take identifiers of procedures as parameters has mode depth 1. For any $k \geq 1$, a procedure that may take identifiers of procedures of mode depth k has mode depth $k + 1$. Thus for Clarke's simulation, procedures of mode depth ≤ 2 suffice. This restriction is also obeyed in PASCAL [JW75]. Procedures of an arbitrary finite mode depth, corresponding to arbitrary *higher types*, as in ALGOL 68 [vWMP⁺75], or even with *self-application*, i.e., with procedure calls of the form *call* $p(\dots, p, \dots)$, as possible in ALGOL 60 [NBB⁺63], are not needed.

6.3. Clarke's language L_4

Clarke also claimed that each language L_i obtained from L_0 by disallowing self-application and the feature $i \in \{1, 2, 3, 4, 5\}$ of the above list has a sound and relatively complete Hoare-like proof system [Cla79]. However, he proved this only for the case of L_3 in which dynamic scope replaces static scope. The completeness argument rests on the fact that each program S in L_3 has a finite range, where the *range* of S is the set of different procedure calls that can be invoked in the computations of S . He claimed that a similar proof system could be obtained for L_4 in which global variables are disallowed. His argument was that programs of L_4 can be transformed into schematically equivalent ones in L_5 , where nested procedure declarations are disallowed. Such programs have a sound and relatively complete Hoare-like proof system.

H. Langmaack and E.R. Olderog [LO80] discovered that this argument is wrong. They considered the *formal execution trees* of programs and showed that programs in L_4 may have trees with context-free path languages whereas programs in L_5 can generate only trees with regular path languages. Therefore they posed the challenge of developing a sound and relatively complete Hoare-like proof system for L_4 .

As a first step, Langmaack proved in [Lan82] that for all ALGOL-like programs in L_4 the divergence problem is decidable for finite interpretations. This is a necessary condition for the existence of a sound and relatively complete Hoare-like proof system according to Lemma A. Moreover, due to a theorem by R.J. Lipton [Lip77], this decidability result is also sufficient for the existence of a sound and relatively complete ‘Hoare logic’ for ALGOL-like programs in L_4 . However, Lipton’s notion of a ‘Hoare logic’ is rather weak: it means that the set of correctness formulas that are true in the sense of partial correctness is recursively enumerable relative to the underlying interpretation. Lipton’s theorem does not yield any usable, syntax-oriented proof rules.

First concrete axiomatizations for L_4 programs appeared independently in papers by Olderog [Old84] and W. Damm and B. Josko [DJ83]. The former studied the case of PASCAL-like programs, i.e. with mode depth ≤ 2 . The key idea is the use of second-order relation variables in the assertion language to stand for the behaviour of uninterpreted, symbolic procedures in a new SEPARATION rule. In order to determine what a procedure call `call $p(\dots q\dots)$` of a procedure p with an actual procedure parameter q does, this rule *first* determines separately what q does and what p does with a symbolic procedure instead of q represented by a formula and *then* composes both results using substitution. Applications of the SEPARATION rule have the effect of a systematic transformation of an initially non-regular formal execution tree into a regular tree at the cost of introducing higher-order tree combinators. The resulting regularity enables a suitably complete Hoare-like proof system.

The paper [DJ83] handled the more general case of ALGOL-like programs with arbitrary finite modes by using higher-order predicate variables and unevaluated substitutions to such variables. Common to both papers is that they deviated from the standard notion of relative completeness in that they used a higher-order assertion language and an appropriate notion of expressiveness.

The status of the language L_4 was finally clarified in 1989 in an over 90 pages long paper by S.M. German, E.M. Clarke and J.Y. Halpern [GCH89] in which a Hoare-like proof system for L_4 was provided that is sound and relatively complete in the (original) sense of Cook. The completeness proof heavily relies on an appropriate arithmetic encoding of procedure declarations so that the assertion language remains first-order.

6.4. The characterization problem

Exploring the border between relative completeness and incompleteness (in the sense of Cook) of Hoare’s logics has been the focus of considerable research in the 1980s. Clarke called it the “Characterization Problem for Hoare Logics” in his survey article [Cla85]. The key question is what is meant by a ‘Hoare logic’. Several researchers proved results for a weak interpretation where a Hoare logic is just the set of correctness formulas $\{P\} S \{Q\}$ that are true in the sense of partial correctness. Here the problem is for which classes of programs S , assertions P , Q , and underlying interpretations I this set is recursively enumerable or even decidable.

We cited already Lipton’s result [Lip77]. Since its proof was only sketched, Langmaack provided in [Lan79] a rigorous proof for the setting of ALGOL-like programs S . Both Lipton and Langmaack restrict themselves to quantifier-free assertions P , Q . Clarke, German and Halpern established [CGH83] an “effective axiomatization of Hoare logics” extending Lipton’s result to first-order formulas P , Q . However, these results do not yield any usable, syntax-directed proof rules that are in the spirit of Hoare-like proof systems that appeared since the publication of [Hoa69]. Therefore Clarke stated in [Cla85]:

“Certainly the most important research problem is to develop a version of the characterization theorem that provides some insight as to when a syntax-directed proof system can be obtained.”

Two such characterization theorems were established in [Old81] and [Old83a]. The first paper studied a language L_{Algol} of ALGOL-like programs with the features 1, 2, 4 and 5 of [Cla79] (see Sect. 6.2). Feature 3 (static scope) was not fixed but left as a parameter in the form of a *copy rule*. Such a rule defines how during program execution a procedure call is replaced by a copy of the procedure body. In the copy certain occurrences of local variables and procedures in the procedure body are renamed to avoid name clashes. By varying the copy rule, different program semantics varying from dynamic to static scope were defined. The COPY rule from Sect. 5.2 is a trivial example of such a rule in which no renaming takes place.

Parameterized by a given copy rule C , a Hoare-like proof system $\mathcal{H}(C)$ for the partial correctness of ALGOL-like programs was introduced. A program S was called *C-bounded* if applications of the copy rule C do not lead to programs with “procedural reference chains” of arbitrary length. A program S has a *finite C-index* if the relation of “substitutional equivalence” induces only finitely many equivalence classes in the set of reachable procedure

calls. Due to this equivalence, the condition of finite \mathcal{C} -index is more liberal than the conditions of *finite recursive cycle* and *finite range* used by Gorelick [Gor75] and Clarke [Cla79], respectively, in their completeness proofs. The following result holds.

Theorem Suppose the assertion language is expressive relative to the considered interpretation I and the copy rule \mathcal{C} . Then the following are equivalent for ALGOL-like programs S and assertions P, Q :

1. $\{P\} S \{Q\}$ can be proved in the proof system $\mathcal{H}(\mathcal{C})$ assuming all true formulas in I .
2. S has a finite \mathcal{C} -index and $\{P\} S \{Q\}$ is true in I in the sense of partial correctness.
3. S is \mathcal{C} -bounded and $\{P\} S \{Q\}$ is true in I in the sense of partial correctness.

As corollaries to this theorem various completeness results were obtained in [Old81]. For the naive copy rule \mathcal{C}_n , which models “dynamic scoping”, the proof system $\mathcal{H}(\mathcal{C}_n)$ is sound and relatively complete for the full set L_{Algol} because all programs in L_{Algol} are \mathcal{C}_n -bounded. This corresponds to Clarke's language L_3 . For the Algol 60 copy rule \mathcal{C}_{60} , which models “static scoping”, the proof system $\mathcal{H}(\mathcal{C}_{60})$ is sound and relatively complete for the following sublanguages of L_{Algol} because all programs in these sublanguages are \mathcal{C}_{60} -bounded:

- L_{pnes} – all programs without procedure nesting, corresponding to Clarke's language L_5 ,
- L_{par} – all programs with parameterless procedures only,
- L_{pp} – all programs without procedures as parameters, corresponding to Clarke's language L_1 ,
- L_{rp} – all programs in which formally recursive procedures are disallowed to have procedure identifiers as formal parameters, a superset of Clarke's language L_2 ,
- L_{gf} – all programs without global formal procedure parameters.

The paper [Old83a] studied a set L_{Pas} of programs with PASCAL-like procedures S , i.e., with mode-depth ≤ 2 . Here a sublanguage $L \subseteq L_{Pas}$ is called *admissible* if it is closed under certain program transformations that leave the procedure structure invariant. In particular, these transformations allow for the introduction of global variables. A tree T is *regular* if the set of paths in T is a regular language in the Chomsky hierarchy or, equivalently, if T has only finitely many different patterns of subtrees. The *formal call tree* of a program S describes in which order the procedures of S are called, where a branching appears when from one procedure several immediate successors can be called. The following theorem refers to a particular Hoare-like proof system H_0 that was introduced in [Old83a].

Theorem For every admissible language $L \subseteq L_{Pas}$ the following are equivalent:

1. There exists a sound and relatively complete Hoare logic for L in the sense of Lipton.
2. The divergence problem for L is decidable for finite interpretations.
3. All programs in L have regular formal call trees.
4. The Hoare-like proof system H_0 is sound and relatively complete for L .

Since L_{Pas} contains programs with a non-regular formal call tree, the theorem implies that there is no sound and relatively complete Hoare logic for L_{Pas} itself, a fact that follows from Clarke's incompleteness result. Note that this theorem does not contradict the results obtained for L_4 because a sublanguage $L \subseteq L_{Pas}$ without global variables is not admissible.

7. Nondeterministic and Probabilistic Programs

7.1. Reasoning about nondeterminism

In the context of programming languages *nondeterminism* stands for the phenomenon that a program can yield more than one answer. In the sixties and early seventies a couple of simple programming constructs were proposed that introduced nondeterminism. In particular, in [Lau71] the nondeterministic statement S_1 or S_2 was considered with the meaning: execute either S_1 or S_2 , and the following proof rule:

OR

$$\frac{\{P\} S_1 \{Q\}, \{P\} S_2 \{Q\}}{\{P\} S_1 \text{ or } S_2 \{Q\}}$$

But the undoubtedly most successful and elegant proposal is the language of *guarded commands* introduced by E.W. Dijkstra in [Dij75] and in the book form in [Dij76a]. Dijkstra's original motivation for introducing this language was to simplify programs by delaying some arbitrary choices to the implementation level and to restore symmetry that is not present in the **if-then-else** statement. This led to a specific proposal for nondeterminism.

For the sake of what follows it is sufficient to consider two crucial statements of Dijkstra's language:

- the *alternative command*

if $B_1 \rightarrow S_1 \square \dots \square B_n \rightarrow S_n$ **fi**,

- the *repetitive command*

do $B_1 \rightarrow S_1 \square \dots \square B_n \rightarrow S_n$ **od**,

where each B_i is a Boolean expression, called a *guard*, and each S_i is a program statement. The symbol \square represents a nondeterministic choice.

The alternative command is executed by selecting a guard B_i that evaluates to **true** and executing the associated statement S_i . If more than one guard B_i evaluates to true any of the corresponding statements S_i may be executed next. If all guards evaluate to **false**, the execution of the alternative command results in a *failure*. The repetitive command is executed in a similar way, with two differences. First, after termination of a selected statement S_i the repetitive command is executed again. Second, if all guards evaluate to **false**, the execution of the repetitive command simply terminates. So it is a natural generalization of the **while** statement.

As an illustration of the use of the alternative command consider the customary program for computing the maximum of two numbers using the conditional statement:

if $x \geq y$ **then** $max := x$ **else** $max := y$ **fi**.

A solution using the alternative command is symmetric in the variables x and y and also involves nondeterminism:

if $x \geq y \rightarrow max := x \square y \geq x \rightarrow max := y$ **fi**.

As an illustration of the use of the repetitive command consider the customary **while** program for computing the *greatest common divisor* (*gcd*) of two natural numbers, initially stored in the variables x and y :

while $x \neq y$ **do**
 if $x > y$ **then** $x := x - y$ **else** $y := y - x$ **fi**
od.

Using the repetitive command the same algorithm can be written as

do $x > y \rightarrow x := x - y \square y > x \rightarrow y := y - x$ **od**.

Both programs terminate with the *gcd* stored in the variables x and y but the second program is more readable and, unlike the first one, is symmetric in the variables x and y .

To reason about the guarded command language Dijkstra introduced in [Dij75] the *weakest precondition calculus* that we briefly discuss in Sect. 10. This approach allows one to reason about total correctness. It is easy to conceive Hoare-style proof rules that deal with partial correctness of the alternative and repetitive commands (they were proposed first in [dB80]). For total correctness some care has to be exercised to deal with the absence of failures, which is a new concept in this framework. The appropriate rule, introduced in [Apt84], takes the following form:

ALTERNATIVE COMMAND II

$$\frac{P \rightarrow \bigvee_{i=1}^n B_i, \quad \{P \wedge B_i\} S_i \{Q\}, i \in \{1, \dots, n\}}{\{P\} \text{ if } \bigwedge_{i=1}^n B_i \rightarrow S_i \text{ fi } \{Q\}}$$

To obtain a proof system for total correctness of guarded commands it suffices now to deal with the termination of the repetitive command. The following natural generalization of the WHILE II rule was used in [AdBO9]:

REPETITIVE COMMAND II

$$\frac{\{P \wedge B_i\} S_i \{P\}, i \in \{1, \dots, n\}, \quad \{P \wedge B_i \wedge t = z\} S_i \{t < z\}, i \in \{1, \dots, n\}, \quad P \rightarrow t \geq 0}{\{P\} \text{ do } \bigwedge_{i=1}^n B_i \rightarrow S_i \text{ od } \{P \wedge \bigwedge_{i=1}^n \neg B_i\}}$$

where t is a termination function and z is an integer variable not occurring in P , t , B_i or S_i for $i \in \{1, \dots, n\}$.

7.2. Reasoning about fairness

Fairness is a concept that arises in presence of any form of nondeterministic choice. Suppose that we repeatedly have some choice among a fixed set of alternatives, for instance of going left or going right. If we repeatedly select the alternative of ‘going left’, then we systematically ignore the other alternative. In such case we can argue that the adopted selection procedure is not fair with respect to the other alternative, ‘going right’. To exclude such unfair selection procedures we need to focus on infinite ‘runs’ of selections of alternatives and make precise when an alternative can be selected.

These matters can be discussed in a more precise way using the guarded commands language. Consider the following program, where k is a fixed natural number:

$$x := 1; \text{ do } x > 0 \rightarrow x := x + 1 \square x > k \rightarrow x := 0 \text{ od.}$$

It does not always terminate, since we can repeatedly select the first guard. The resulting computation is considered *unfair* since from some moment on the second guard is always enabled (i.e., evaluates to **true**), but never selected. Once the second guard is selected when it is enabled, the program terminates. More formally, we say that the above program terminates under the *fairness* assumption, which stipulates that each guard that is from some moment on continuously enabled is infinitely often selected.

As another example consider the following program, where $odd(x)$ is a test with the expected meaning:

$$x := 1; \text{ do } x > 0 \rightarrow x := x + 1 \square odd(x) \rightarrow x := 0 \text{ od.}$$

Also this program does not always terminate. The only infinite computation repeatedly ignores the second guard. However, in contrast to the previous example, in this infinite computation at no moment the second guard becomes continuously enabled. So even under the fairness assumption this program does not terminate.

On the other hand in this infinite computation the second guard is infinitely often enabled. We say that this computation is *strongly unfair*. If the second guard is selected when it is enabled, the program terminates. In this case we say that the above program terminates under the *strong fairness* assumption, which stipulates that each guard that is infinitely often enabled is infinitely often selected. (To stress the difference the first notion of fairness is usually called *weak fairness*.)

Both forms of fair termination can be established by means of appropriate proof rules. In what follows we explain the *transformational approach* proposed in [AO81] and in a journal form in [AO83]. To prove termination of a repetitive command S under the assumption of weak (or strong) fairness (with respect to a precondition P) one can transform it into a repetitive command $T(S)$ the computations of which coincide with the weakly (or strongly) fair computations of S . To this end, we need a *random assignment* $x := ?$ that assigns nondeterministically to the variable x an arbitrary natural number. To understand why this command naturally arises when considering fairness, note that under both assumptions of fairness the program

$$b := \text{true}; x := 0; \text{ do } b \rightarrow x := x + 1 \square b \rightarrow b := \text{false} \text{ od}$$

mentioned in [Dij76a] is equivalent to

$$b := \text{false}; \quad x := ?.$$

In what follows we limit ourselves to the presentation of strong fairness and to explain the idea we focus on a repetitive command with just two guards:

$$S \equiv \mathbf{do} \ B_1 \rightarrow S_1 \ \square \ B_2 \rightarrow S_2 \ \mathbf{od}.$$

The transformed program $T(S)$ uses two auxiliary variables z_1 and z_2 ranging over natural numbers and has the following form:

$$\begin{aligned} T(S) \equiv & z_1 := ?; \quad z_2 := ?; \\ & \mathbf{do} \\ & \quad B_1 \wedge z_1 \leq z_2 \rightarrow S_1; \quad z_1 := ?; \quad \mathbf{if} \ B_2 \ \mathbf{then} \ z_2 := z_2 - 1 \ \mathbf{fi} \\ & \quad \square \\ & \quad B_2 \wedge z_2 \leq z_1 \rightarrow S_2; \quad z_2 := ?; \quad \mathbf{if} \ B_1 \ \mathbf{then} \ z_1 := z_1 - 1 \ \mathbf{fi} \\ & \mathbf{od}. \end{aligned}$$

Informally, each variable z_i tracks the number of times the corresponding guard is enabled and the augmented guards prevent that an infinitely often enabled guard is never selected.

To reason about the strong fair termination of S it is now equivalent to reason about the termination of the program $T(S)$. To this end, we only need an axiom dealing with the random assignment. Such an axiom was proposed in [Har79]:

RANDOM ASSIGNMENT

$$\{P\} \ x := ? \ \{P\}$$

where x does not appear free in P .

But this indirect approach can be avoided by absorbing the transformation $T(\cdot)$ into the proof of termination of the program $T(S)$. This way one obtains a proof rule for establishing termination under the strong fairness assumption that deals directly with the original program S and is similar in shape to the REPETITIVE COMMAND II rule. We omit the details, though we should mention that to reason about fairness it is in general necessary that the termination function takes values from an arbitrary well-founded ordering and not just natural numbers (see [AO83], [APS84], and [AP86], where soundness and relative completeness of a proof system for total correctness of **while** programs augmented with random assignment was established.).

Different approaches to reason about fairness in the context of nondeterministic programs were independently proposed in [LPS81] and [GFMDR81], that appeared in a journal form as [GFMDR85]. In [APS84] the method presented here was extended to programs with nested nondeterminism. These approaches and proof rules were discussed in book form in [Fra86], where also additional versions of fairness were considered.

The transformational approach to fairness was originally developed in [AO81, AO83] for the Dijkstra's guarded command language in which each repetitive command has a fixed finite number of alternatives. The transformations can be seen as implementing a general fair scheduler controlling finitely many processes. These transformations were extended in [OP10, HOP10] to deal with *dynamic control*, where processes can be created dynamically. Then the overall number of processes can be infinite, but at each step of an execution of the system the number of created processes is finite. J. Hoenicke and A. Podelski went in [HP15] one step further and extended the transformations to deal with fairness with an *infinitary control*, where the number of created processes can be infinite. Both dynamic and infinitary control were expressed by repetitive commands with infinitely many alternatives. However, these papers did not propose any proof rules derived from the new transformations.

7.3. Probabilistic programs

Probabilistic programs are sequential programs with the ability to draw values at random from probabilistic distributions. They have attracted large attention in the research community due to many applications, for example in security to describe randomized encryptions, in machine learning to describe distribution functions, and in randomized algorithms. They have typically just a few lines of code, but are hard to analyze, see, e.g., [KGJ⁺15]. Properties of interest for such programs include the *expected runtime* and *almost sure termination*.

Most formal modelling takes place in the setting of an extended Dijkstra's guarded command language, called *probabilistic guarded command language*, abbreviated pGCL. Here, both nondeterministic choice and probabilistic choice are admitted. In particular, McIver and Morgan [MM05] carried out their research on probabilistic programs in this setting. They extended the notion of weakest precondition to *weakest pre-expectations*.

J. den Hartog and E.P. de Vink [dHdV02] introduced a Hoare-like proof system for partial correctness of probabilistic programs which are defined as (deterministic) **while** programs extended by the *probabilistic choice* $S_1 \oplus_\rho S_2$ between the statements S_1 and S_2 . This intention is that the statement S_1 is chosen with the probability ρ and the statement S_2 with the probability $1 - \rho$. This necessitates the introduction of probabilistic predicates in the assertion language. In these predicates, the real-valued expression $\mathbb{P}(R)$ yields the probability that the normal predicate R holds. For example, for an integer variable x the correctness formula

$$\{x = 1\} x := x + 1 \oplus_{\frac{1}{3}} x := x + 2 \{ \mathbb{P}(x = 2) = \frac{1}{3} \wedge \mathbb{P}(x = 3) = \frac{2}{3} \}$$

holds. The proof system of [dHdV02] contains several new rules that go beyond the ones concerning the **while** programs. In particular, for the probabilistic choice the following rule was proposed:

PROBABILISTIC CHOICE

$$\frac{\{P\} S_1 \{Q\}, \{P\} S_2 \{Q'\}}{\{P\} S_1 \oplus_\rho S_2 \{Q \oplus_\rho Q'\}}$$

where the probabilistic choice operator \oplus_ρ is also applied to the probabilistic predicates Q and Q' .

Den Hartog and de Vink established soundness of their proof system. They also proved relative completeness for the subset of loop-free probabilistic programs and a restricted set of predicates in the postcondition. The proof is based on calculating the weakest precondition.

Some other developments concerning verification of probabilistic programs are discussed in Sect. 10.7.

8. Parallel and Distributed Programs

By a *concurrent program* we mean a program that has a number of components, the execution of which proceeds in parallel. If the program components share some variables one usually refers to this framework as shared memory parallelism. Following [AdBO09], we refer to the corresponding programs as *parallel programs*. If the program components do not share variables but can communicate by messages, one usually calls these components *processes* and refers to the corresponding concurrent programs as *distributed programs*. A most commonly used approach relies on asynchronous message passing. However, within the framework of program verification, for reasons of simplicity, one rather focuses on synchronous message passing, which means that the processes need to synchronize their actions to make the communication possible. Again following [AdBO09], we mean by *distributed programs* the ones with synchronous message passing, within the context of a specific language proposal due to Hoare [Hoa78].

8.1. Reasoning about parallel programs

Verification of parallel programs calls for new insights. To discuss the matters let us first introduce the syntax. By

$$[S_1 \parallel \dots \parallel S_n]$$

we mean a *parallel composition* of n sequential programs, S_1, \dots, S_n , that may share variables.

The main complication in reasoning about parallel programs is the interference caused by the use of shared

variables. For example, both

$$\{x = 0\} [x := x + 1; x := x + 1] \{x = 2\}$$

and

$$\{x = 0\} x := 2 \{x = 2\}$$

hold but

$$\{x = 0\} [x := x + 1; x := x + 1 \parallel x := 2] \{x = 2 \vee x = 4\}$$

does not since one possible computation consists of executing the assignments in the following order $x := x + 1$; $x := 2$; $x := x + 1$, which yields the final value 3.

The first limited approach to the verification of parallel programs was proposed by Hoare in [Hoa72c], where he dealt with parallel composition of disjoint program components. More formally, we say that the component programs S_1, \dots, S_n of $[S_1 \parallel \dots \parallel S_n]$ are *disjoint* if no variable subject to change in one component appears in another component. This led to the following proof rule:

DISJOINT PARALLELISM

$$\frac{\{P_i\} S_i \{Q_i\}, i \in \{1, \dots, n\}}{\{\bigwedge_{i=1}^n P_i\} [S_1 \parallel \dots \parallel S_n] \{\bigwedge_{i=1}^n Q_i\}}$$

where $free(P_i, Q_i) \cap change(S_j) = \emptyset$ for $i \neq j$.

As noticed by Hoare in [Hoa75], parallel and sequential compositions of program components coincide under the assumption of disjointness. This brought him to suggest the following rule:

SEQUENTIALIZATION

$$\frac{\{P\} S_1; \dots; S_n \{Q\}}{\{P\} [S_1 \parallel \dots \parallel S_n] \{Q\}}$$

(In [Hoa72c] and [Hoa75] actually only parallel composition of two components was studied). In 1975 this limited approach to program correctness was extended to arbitrary parallel programs. It was based on a novel idea of *interference freedom*, first introduced in the PhD thesis of S. Owicki [Owi75], and subsequently published in two papers by her and supervisor D. Gries, [OG76a] and [OG76b]. This approach became known as the *Owicki-Gries* method.

In what follows we explain the ideas behind this extension of Hoare's logic. To deal with the problem of interference mentioned above, Owicki and Gries suggested to compose not the correctness statements about the component programs *but* their proofs, presented in an appropriate form. This form relies on the fact that the correctness proofs of the component programs are syntax-directed, that is they follow in some sense the program structure. As a result the proof can be recorded by retaining the used assertions in the program text. For example, the application of the WHILE rule can be retained in the text of the **while** statement by writing its conclusion as

$$\{P\} \mathbf{while} B \mathbf{do} \{P \wedge B\} S \{P\} \mathbf{od} \{P \wedge \neg B\},$$

while the conclusion of the CONSEQUENCE rule can be written as

$$\{P\} \{P_1\} S \{Q_1\} \{Q\},$$

which amounts to interpreting two consecutive assertions as an implication. The crucial point of such a proof representation, called a *proof outline*, is that each statement S is preceded in it by at least one assertion. Subsequent assertions indicate implication between them.

Here is an example of a proof outline for the correctness proof of the DIV program that we considered in Sect. 3.1:

```

{true}
{x = x + y · 0}
r := x;
{x = r + y · 0}
q := 0;
{P}
while y ≤ r do
  {P ∧ y ≤ r}
  {x = (r - y) + y · (1 + q)}
  r := r - y;
  {x = r + y · (1 + q)}
  q := 1 + q
  {P}
od
{¬y ≤ r ∧ P},

```

where

$$P \equiv x = r + y \cdot q.$$

In the Appendix we give another example of a proof outline by representing in such a form Turing's proof discussed in Sect. 2.1. To understand the essence of the approach of Owicki and Gries let us return for a moment to the original Hoare's proof system. It is sound in the following stronger sense, where, given a proof outline and a subprogram T , we denote by $pre(T)$ the assertion that directly precedes T . For a proof see, e.g., [AdBO09].

Strong Soundness Theorem Consider a proof outline that represents the proof of a correctness formula $\{P\} S \{Q\}$. Take a computation ξ of S that starts in a state that satisfies P . Each time ξ reaches a substatement T of S , the precondition $pre(T)$ is satisfied. Further, if ξ terminates, its final state satisfies Q .

Armed with this knowledge let us return to parallel programs. Suppose that we established the correctness formulas $\{P_i\} S_i \{Q_i\}$, where $i \in \{1, \dots, n\}$ and S_1, \dots, S_n are the component programs of the parallel program $[S_1 \parallel \dots \parallel S_n]$. Let $\{P_i\} S_i^* \{Q_i\}$ be the corresponding proof outlines. We call them *interference-free* if for all assignments $x := s$ in $[S_1 \parallel \dots \parallel S_n]$ and all assertions R used in a proof outline of another component the correctness formula

$$\{R \wedge pre(x := s)\} x := s \{R\}$$

holds. Informally, the proof outlines are interference-free if the execution of each assignment statement in the state that satisfies its assertion does not invalidate the assertions used in the proof outlines of other components. Then the following proof rule allows us to reason about parallel programs:

PARALLELISM

The proof outlines $\{P_i\} S_i^* \{Q_i\}$, $i \in \{1, \dots, n\}$, are interference-free

$$\frac{}{\{\bigwedge_{i=1}^n P_i\} [S_1 \parallel \dots \parallel S_n] \{\bigwedge_{i=1}^n Q_i\}}$$

Owicki and Gries noted that to reason about parallel programs they needed *auxiliary variables* (occasionally called by other researchers *ghost variables*), so variables that neither influence the control flow nor the data flow of the program, but only record some additional information about the program execution. Formally, a set of simple variables A is called a *set of auxiliary variables* of S if each variable from A occurs in S only in assignments to the variables from A . The appropriate proof rule allowing one to delete them is:

AUXILIARY VARIABLES

$$\frac{\{P\} S \{Q\}}{\{P\} S_A \{Q\}}$$

where for some set of auxiliary variables A of S the program S_A results from S by deleting all assignments to the variables in A , and no variable from A appears free in Q .

In presence of shared variables it is sometimes useful to restrict the number of interference points. This can be done by allowing *atomic sections* (see [Lip75] and [Lam77]), which are statements of the form $\langle S \rangle$, where S does not contain any loops or further atomic sections. The idea is that such a statement is executed without any interruption by another component of the parallel program. This suggests the following proof rule (see [Owi78]):

ATOMIC SECTION

$$\frac{\{p\} S \{q\}}{\{p\} \langle S \rangle \{q\}}$$

To allow for synchronization between the component programs, Owicki and Gries used a generalization of the atomic section statement, called an *await statement* **await** B **then** S **end**, with the following meaning: if B evaluates to true, the statement S is executed as an atomic section. The corresponding proof rule is:

AWAIT

$$\frac{\{P \wedge B\} S \{Q\}}{\{P\} \text{await } B \text{ then } S \text{ end } \{Q\}}$$

The presence of synchronization statements leads to a possibility of a *deadlock*, an undesired situation in which some components of a parallel program did not terminate while the nonterminated components are all blocked. Owicki and Gries noted that interference-free outlines allow one to reason about absence of deadlock in a simple way. To this end, it suffices to identify the possible deadlock situations and to show for each of them that the corresponding conjunction of the $pre(T)$ assertions cannot be simultaneously satisfied.

This approach to verification was successfully applied in [OG76a] to establish correctness of non-trivial parallel programs, in particular two versions of the classical producer/consumer problem, a parallel version of the FIND program mentioned in Sect. 3.1, and an implementation of Dijkstra's semaphores. In [OG76b] it was extended to parallel programs that use a more efficient synchronization construct called conditional regions introduced in [Hoa72a], and in [Owi78] to concurrent programs with shared data classes.

Termination Let us discuss now the issue of termination of parallel programs. Owicki and Gries proposed to establish it by using the WHILE II rule instead of the WHILE rule and by postulating that no component program increases the termination function of a **while** statement used in another component program. However, the authors of [AdBO90] noticed that the termination of parallel programs is subtler than it sounds and came up with an example of simple parallel program that does not terminate, even though its termination can be established using Owicki and Gries method.

The problem has to do with the fact that the assertions used in the proof of the second premise of the WHILE II rule, so $\{P \wedge B \wedge t = z\} S \{t < z\}$, are not tested for interference freedom. In [AdBO90] two ways of solving this problem are proposed. The first is to retain in the proof outlines the correctness proofs of both the first and second premise of this rule, so that all assertions used in these proofs are subjected to the interference freedom test. Another is by adding additional requirements to the definition of interference freedom that ensure that the used termination functions decrease along all syntactically possible paths through the program. The latter approach suffices to validate the termination proofs presented in [OG76a].

Finally, termination of parallel programs under the fairness assumption was dealt with in [OA88], where the Owicki-Gries method was combined with a transformational approach presented in Sect. 7.2.

Independently of the Owicki-Gries method of [OG76a], L. Lamport proposed in [Lam77] an essentially equivalent approach to verification of parallel programs. The difference was that, as in [Flo67], he considered programs presented as flowcharts, now one for each component program. As a result the interference-freedom test referred to the assertions attached to the flowchart nodes and translated into a requirement that these assertions are *monotone*, that is, that they are maintained by the actions of the other components. Lamport used this approach to establish correctness of a solution to the mutual exclusion problem. In the paper also program properties were considered that had not been hitherto addressed. We shall return to this matter in the last section of the paper.

The rely-guarantee approach A drawback of the Owicki-Gries method is that due to the test of interference freedom, verification of a parallel program $S \equiv [S_1 \parallel \dots \parallel S_n]$ is possible only if all components S_1, \dots, S_n of S are explicitly given. Further, this method is not compositional, where 'compositionality' means that (partial) correctness of a parallel program is derived directly from the correctness of its components.

This prompted research on alternative formalisms for reasoning about parallel programs. A discussion of these approaches can be found in [dRdBH⁺01, pp. 479-484]. Here we limit ourselves to an account of one, perhaps most successful proposal, called the *rely-guarantee approach*. It was introduced in the PhD thesis of C.B. Jones, [Jon81], the essence of which appeared in [Jon83]. This approach provides a compositional way reasoning about concurrent programs by incorporating the interference-freedom test into the proof. This is achieved by using more informative correctness formulas.

We follow here the presentation given in the book by W.P. de Roever et al. [dRdBH⁺01]. In the rely-guarantee approach one assumes that a given program S is executed in some environment and therefore uses an extended specification format

$$\langle R, G \rangle : \{P\} S \{Q\},$$

where a correctness formula $\{P\} S \{Q\}$ is extended by an *interface specification* $\langle R, G \rangle$ of the environment consisting of a rely condition R and a guarantee condition G . In contrast to the assertions P and Q in the correctness formula, R and G are *predicates on transitions*, i.e., they relate two states, the one before executing a transition and the one after it. As in Turing's flowchart in Sect. 2.1, primed versions of variables are used to refer to the state after executing the transition. For example, $x' < x$ expresses that the value of the variable x decreases. The idea is that R states assumptions that S makes on the transitions of its environment and that G states the guarantees that S provides to the environment.

Informally, $\langle R, G \rangle : \{P\} S \{Q\}$ expresses that the correctness formula $\{P\} S \{Q\}$ is true in the sense of partial correctness if

- whenever at some moment during the computation of S all past environmental transitions satisfy R
- then all transitions of S up to that moment satisfy G .

For parallel composition of programs S_1 and S_2 the following proof rule is presented in [dRdBH⁺01]:

PARALLEL COMPOSITION

$$\frac{\begin{array}{l} (R \vee G_1) \rightarrow R_2, \\ (R \vee G_2) \rightarrow R_1, \\ (G_1 \vee G_2) \rightarrow G, \\ \langle R_i, G_i \rangle : \{P_i\} S_i \{Q_i\}, i = 1, 2, \end{array}}{\langle R, G \rangle : \{P_1 \wedge P_2\} [S_1 \parallel S_2] \{Q_1 \wedge Q_2\}}$$

In [dRdBH⁺01, p. 453] this rule is explained as follows:

- every transition of S_i (which is characterized by G_i) and every transition of the common environment of S_1 and S_2 (characterized by R) is seen by S_j with $i \neq j$ as an environment transition which has to satisfy R_j ,
- every transition by S_1 or S_2 is a transition of $[S_1 \parallel S_2]$, and therefore has to satisfy G , and
- since the validity of $\langle R_i, G_i \rangle : \{P_i\} S_i \{Q_i\}$ implies Q_i is invariant under R_i , the postcondition $Q_1 \wedge Q_2$ holds upon termination of $[S_1 \parallel S_2]$; since R implies R_i , $i = 1, 2$, this implies that $Q_1 \wedge Q_2$ is invariant under R after $[S_1 \parallel S_2]$ has terminated, too.

Note that the rule is compositional for the customary correctness formulas. The crux of applying it is to find suitable rely-guarantee conditions for which the implications in the premises hold. In [Jon83] and [dRdBH⁺01] this approach was illustrated by providing alternative correctness proofs of a solution to the mutual exclusion problem and of a parallel version of the FIND program considered in [OG76a]. In practice, this approach may be as difficult as proving interference freedom.

8.2. Reasoning about distributed programs

Hoare introduced an elegant approach to distributed programming based on synchronous communication [Hoa78]. In an intentional analogy to the title of Dijkstra's seminal paper [Dij68] on parallel programs, this proposal was called Communicating Sequential Processes. In the paper Hoare introduced a simple programming language for distributed programming, called since then CSP, in which Dijkstra's guarded command language was extended by allowing communication primitives for synchronous communication. This focus on synchronous, as opposed to asynchronous, communication had a huge impact on the theory of distributed programming and was also realized in the programming language Occam [INM84]. The idea is that synchronous communication is simpler to reason about as it obviates the discussion of message ordering and buffers and their management. Synchronous

communication can be implemented by means of asynchronous one, by adding additional processes that simulate message buffers, so it can be viewed as an elegant abstraction.

In [AFdR80] a proof system was proposed to reason about a simple class of CSP programs. The crucial idea of this approach was an introduction of a *cooperation test* that corresponds to the interference test of the Owicki-Gries method.

To explain the matters we first clarify the relevant aspects of CSP. Each CSP program consists of a parallel composition of processes, written as

$$[PR_1 :: S_1 \parallel \dots \parallel PR_n :: S_n]$$

Each PR_i is a label of a process and S_i is its program. These processes share no variables. They communicate by means of synchronous communication that is achieved by means of two matching *input/output commands*, in short i/o commands. An i/o command has the form $PR_i?x$ (an input command) or $PR_i!t$ (an output command), where x is a variable, and t an expression. The i/o commands $PR_i?x$ or $PR_j!t$ *match* if $i \neq j$, $PR_i?x$ appears in the program for process PR_j , $PR_j!t$ appears in the program for process PR_i , and the types of x and t coincide.

When the control in the programs for processes PR_i and PR_j is just in front of the mentioned i/o commands and they match, they can be executed jointly, with the effect that the value of t is assigned to x . So the effect of the joint execution of the commands $PR_i?x$ or $PR_j!t$ is that of an assignment $x := t$.

In the CSP language, Dijkstra's guarded commands are generalized by allowing i/o commands to appear in the guards. So the guards can now also be of the form $B; \alpha$, where B is a Boolean expression and α is an i/o command. If B evaluates to **true** the i/o command α of the generalized guard behaves the same way as the usual i/o command, though it fails if it addresses a process that terminated. If B evaluates to **false** the generalized guard fails.

Further, Dijkstra's **do-od** notation is replaced by using a star '*' and the '[' and ']' brackets, while the **if-fi** notation is replaced by the '[' and ']' brackets, that are also used to enclose the parallel composition.

As an example of a CSP program consider the following transmission problem, taken from [AdBO9], that is a simplified version of a similar problem discussed in [Hoa78]. We wish to transmit from the *SENDER* process to the *RECEIVER* process through a *FILTER* process a sequence of characters in such a way that *FILTER* process deletes from the sequence all blank characters. The following CSP program

$$[SENDER :: S_1 \parallel FILTER :: S_2 \parallel RECEIVER :: S_3],$$

where

$$\begin{aligned} S_1 &\equiv i := 0; *[i \neq M; FILTER!a[i] \rightarrow i := i + 1], \\ S_2 &\equiv in := 0; out := 0; x := ' '; \\ &\quad *[x \neq '*'; SENDER?x \rightarrow \\ &\quad \quad \text{if } x = ' ' \rightarrow skip \\ &\quad \quad \square x \neq ' ' \rightarrow b[in] := x; \\ &\quad \quad \quad in := in + 1 \\ &\quad \text{fi} \\ &\quad \square out \neq in; RECEIVER!b[out] \rightarrow out := out + 1 \\ &\quad], \\ S_3 &\equiv j := 0; y := ' '; \\ &\quad *[y \neq '*'; FILTER?y \rightarrow c[j] := y; j := j + 1]. \end{aligned}$$

is a solution to this problem.

Here the sequence of characters is initially stored in the array $a[0 : M - 1]$ of characters in the process *SENDER*. The last element of the array is the special character '*', i.e., $a[M - 1] = '*'$. The process *FILTER* has an array b of characters serving as an intermediate store for processing the character sequence and the process *RECEIVER* has an array c of characters to store the result of the filtering process. For coordinating its activities the process *FILTER* uses two integer variables in and out pointing to elements in the array b .

The process *FILTER* can communicate with both other processes. It can receive characters from process *SENDER* until '*' has been received and it can transmit all nonblank characters to the process *RECEIVER*. The Boolean parts of the generalized guards of the *FILTER* process can both evaluate to **true**. In that case the next action can be either a communication between *SENDER* and *FILTER* or between *FILTER* and *RECEIVER*. Consequently this CSP program exhibits a nondeterministic behaviour.

The process *SENDER* terminates once it has sent all its M characters to the *FILTER* process. The process *FILTER* terminates when it has received the character '*' and it has transmitted to *RECEIVER* all nonblank characters it has received. Finally, the process *RECEIVER* terminates once it has received from *FILTER* the character '*'. Thus the parallel composition of these three processes terminates if *SENDER* sends as the last of its M characters the '*'.

The proof system proposed in [AFdR80] extends the one for the guarded commands language by the following axioms and proofs rules that deal with the communication. For simplicity we assume that all variables and expressions are of the same type.

INPUT

$$\{P\} PR_i?x \{Q\}$$

OUTPUT

$$\{P\} PR_i!t \{P\}$$

GENERALIZED ALTERNATIVE COMMAND

$$\frac{\{P \wedge B_i\} \alpha_i \{R_i\}, \{R_i\} S_i \{Q\}, i \in \{1, \dots, n\}}{\{P\} [\square_{i=1}^n B_i; \alpha_i \rightarrow S_i] \{Q\}}$$

GENERALIZED REPETITIVE COMMAND

$$\frac{\{P \wedge B_i\} \alpha_i \{R_i\}, \{R_i\} S_i \{P\}, i \in \{1, \dots, n\}}{\{P\} *[\square_{i=1}^n B_i; \alpha_i \rightarrow S_i] \{P \wedge \bigwedge_{i=1}^n \neg B_i\}}$$

The INPUT axiom may look strange since it allows us to conclude an arbitrary postcondition. However, the used assertions still have to pass the cooperation test. This test refers to the proof outlines which are defined as in the case of Owicki-Gries approach.

Suppose now that we established the proof outlines $\{P_i\} S_i^* \{Q_i\}$, where $i \in \{1, \dots, n\}$ and S_1, \dots, S_n are respective programs of the processes PR_1, \dots, PR_n . We say that these proof outlines *cooperate* if

- the assertions used in $\{P_i\} S_i^* \{Q_i\}$ contain no variables subject to change in S_j for $i \neq j$,
- $\{pre_1 \wedge pre_2\} PR_j?x \parallel PR_i!t \{post_1 \wedge post_2\}$ holds whenever $\{pre_1\} PR_j?x \{post_1\}$ and $\{pre_2\} PR_i!t \{post_2\}$ are taken respectively from the proof outlines $\{P_i\} S_i^* \{Q_i\}$ and $\{P_j\} S_j^* \{Q_j\}$.

Intuitively, proof outlines cooperate if they help each other to validate the post conditions of the i/o commands present in these proofs. To establish cooperation the following axiom is needed.

COMMUNICATION

$$\{\text{true}\} PR_j?x \parallel PR_i!t \{x = t\}$$

provided $PR_j?x$ and $PR_i!t$ match.

This axiom simply states that, as mentioned before, the effect of the joint execution of a pair of matching i/o commands $PR_j?x$ and $PR_i!t$ is that of an assignment $x := t$. Recall that in CSP, the processes composed by parallel composition do not share variables, so x cannot occur in t . This justifies the simple postcondition $x = t$ for the assignment $x := t$.

Then the following proof rule allows one to draw conclusion about the parallel composition of processes:

CSP PARALLELISM

The proof outlines $\{P_i\} S_i^* \{Q_i\}$, $i \in \{1, \dots, n\}$, for the processes PR_1, \dots, PR_n cooperate

$$\frac{}{\{\bigwedge_{i=1}^n P_i\} [PR_1 :: S_1 \parallel \dots \parallel PR_n :: S_n] \{\bigwedge_{i=1}^n Q_i\}}$$

To reason about the CSP programs, as in the case of the Owicki-Gries approach, the AUXILIARY VARIABLES rule is needed. The presented reasoning about deadlock freedom is analogous as in [OG76a], though the

reasons for a deadlock can now be different. In particular, a process can be blocked forever if the control in its program is just before an i/o command that addresses a process that terminated. Using the proposed proof system some example CSP programs were proved correct in [AFdR80]. (In [Moi83] it was pointed out that one of the correctness proofs contained an error and a corrected version was presented.) The proof system was subsequently proved in [Apt83] to be sound and complete in the sense of Cook. Neither [AFdR80] nor [Apt83] considered termination.

Independently of [AFdR80] a very similar proof system to reason about CSP programs was proposed in [LG81], in which the *satisfaction property* corresponds to the cooperation test. However, the authors used a different semantics of the generalized repetitive commands than the one stipulated in [Hoa78] and taken care of in the proof system of [AFdR80]. On the other hand, in contrast to [AFdR80], program termination was considered. A number of different approaches to reason about communicating processes was proposed in the literature around that time. They are surveyed in [Hdr86].

Subsequently, a simplified proof system for a fragment of CSP was proposed in [Apt86] and used to establish correctness of a solution to the so-called *distributed termination problem*. In this fragment a program for each process consists of a guarded commands program followed by a single generalized repetitive command. The i/o commands can appear only in the guards of this repetitive command and thus not as separate statements. An example of a CSP program written in this fragment is the above-mentioned solution to the transmission problem. This fragment was studied independently in [ABC87] and [Zöb88], where it was shown that each CSP program can be transformed into a program in this subset using some control variables.

The underlying idea of this approach is that such simpler CSP programs can be transformed into Dijkstra's guarded commands language without introducing any additional variables. By absorbing this transformation into a proof rule one obtains a proof rule that uses a global invariant and deals directly with the considered CSP program. This way this approach dispenses with the cooperation test. Termination is naturally dealt with by following the approach used for the guarded commands language. This proof system was adopted in [AdBO09] and its two previous editions, where its soundness was shown to be a direct consequence of the correctness of the above-mentioned program transformation.

In [AFK88] it was explained that in the context of CSP programs fairness, a notion we discussed in Sect. 7.2, can have various interpretations. One of them states that every pair of matching i/o commands that is infinitely enabled is also infinitely often selected. In [GFK84] a proof rule for dealing with this form of fairness of the CSP programs was proposed.

The approach of [AFdR80] and [LG81] was presented only for CSP programs without nested parallelism. It was extended in [AdB90b] to CSP programs that allowed nested parallelism and also dynamic process creation, a feature not present in CSP. Subsequent work in this direction, [dB91], shifted emphasis to reasoning about objects, a subject that deserves a separate section.

Finally, let us mention that verification of parallel and distributed programs in the Hoare-like style was systematically presented in book form in [dRdBH⁺01].

9. Object-oriented Programs

Object-oriented programming, as exemplified by languages like C++ or Java, builds upon the notion of an object and concepts like inheritance and subtyping. The difficulty in reasoning about such programs is that their execution creates dynamic pointer structures that go beyond the static program structure that has been the backbone of the syntax-directed Hoare-style proof rules discussed so far.

9.1. Language characteristics

As object-oriented programming has been realized in many, often incompatible, ways, we clarify first the main characteristics of the objects here considered. These are:

- objects possess (and *encapsulate*) their own *instance* variables,
- objects interact via *method* calls,
- objects can be dynamically *created*.

Each object consists of a set of instance variables and a set of methods.




Fig. 5. A linked list. This drawing is taken from [AdBO09, p. 190]

In contrast to the formal parameters of procedures and the local variables of the block statements which only exist *temporarily*, the instance variables of an object exist *permanently*. The *local state* of an object is a mapping that assigns values to its instance variables. Each object represents its local state by a *pointer* to it. *Encapsulation* means that the instance variables of an object cannot be directly accessed from other objects; they can be accessed only by the method calls of the object.

A method call invokes a procedure which is executed by the called object. The execution of a method call thus involves a temporary *transfer* of control from the local state of the caller object to that of the called object (also referred to as *callee*). Upon termination of the method call the control returns to the local state of the caller. The method calls are the *only way* of transferring control from one object to another.

The account of verification of object-oriented programs that follows is based on [AdBO09, Chapter 6]. We distinguish two kinds of variables: the set *Var* of *normal* variables, the ones considered so far, and the set *IVar* of instance variables, which are owned by objects.

We consider a set of *methods*, each defined by means of a declaration

$$m(\mathbf{u}) :: S,$$

where the identifier m denotes a method, \mathbf{u} is the list of formal parameters of type *Var*, and S is the *method body*, which may include recursive calls of m .

Methods are invoked by means of the *parameterized method calls* that are of the form

$$s.m(\mathbf{t}).$$

where s is an expression that denotes the *called object* and \mathbf{t} is the list of actual parameters of the method m .

A program consists of a main statement and a set of method definitions. In programs we use a basic (i.e., not compound) type **object** that denotes an infinite set of objects. The constant **null** of type **object** represents the *void reference*, a special construct which does not have a local state. The normal variable **self** of type **object** stores at any moment the currently executing object. Inside a method body no assignments to the variable **self** are allowed, that is, this variable is read-only. Values of type **object** can only be tested for *equality*. Variables of type **object** are called *object variables*, and expressions of type **object** are called *object expressions*.

To illustrate the considered programs consider a recursive method used to find a zero in a linked list represented by objects. We represent such lists using the instance object variable *next* that links the objects of the list, and the constant **null** that allows us to identify the last element of the list. We assume that each object stores a value kept in an instance integer variable *val*. Additionally, we use the normal object variable *first* to point to the first object in the list. Figure 5 shows an example of such a list representation.

The desired method is declared as follows.

```

find :: if val = 0
      then return := self
      else if next ≠ null
            then next.find
            else return := null
      fi
fi
  
```

Then upon termination of the call **self.find** the object variable *return* points to the first object in the list that stores zero, if it exists, and otherwise it returns the void reference, represented by the constant **null**.

9.2. Reasoning about object-oriented programs

The methods in object-oriented programs use local updates to manipulate a global, dynamically changing pointer structure, the *heap*. Therefore the assertion language must be able to express properties of the heap. To this end, *global expressions* are introduced by extending the (local) expressions that may appear inside methods by *navigation expressions* of the form $e.u$, where e is an object expression and u an instance variable. If e denotes a certain object (for example **self**) and u is an instance variable (for example *next*), then $e.u$ points to the corresponding object (in the example **self.next**). This way, one can navigate from object to object along the pointers in the heap. Assertions are then constructed from global Boolean expressions by allowing Boolean combinations and quantification over normal variables in *Var*.

As an example consider the assertion used in [AdBO09, pp. 226-227] to reason about the above method *find*:

$$P \equiv \mathbf{self} = a[k] \wedge a[n] = \mathit{return} \wedge \forall i \in [k : n - 1] : (a[i] \neq \mathbf{null} \wedge a[i].\mathit{val} \neq 0 \wedge a[i].\mathit{next} = a[i + 1]).$$

In addition to the instance object variables *next* and *return* and the instance integer variable *val* used in the definition of *find*, one uses here the following normal variables: an array a of type $\mathbf{integer} \rightarrow \mathbf{object}$, integer variables i, k and n , and the object variable **self**.

The assertion P states that the array section $a[k : n]$ stores a linked list of objects which starts with the object **self**, ends with the object *return*, and all of its objects, except possibly the last one, are different from **null** and do not store in *val* zero. Note the use of the navigation expressions $a[i].\mathit{val}$ and $a[i].\mathit{next}$.

The desired behaviour of the above method *find* can then be specified by means of the following correctness formula:

$$\{\mathbf{true}\} \mathbf{self}.find \{Q\}, \tag{4}$$

where the postcondition Q is defined in terms of the assertion P :

$$Q \equiv (\mathit{return} = \mathbf{null} \vee \mathit{return}.val = 0) \wedge \exists a : \exists k : \exists n \geq k : P.$$

So the postcondition states that the returned object is **null** or stores zero and that for some array section $a[k : n]$ the assertion P holds.

To reason about such correctness formulas we need new axioms and proof rules.

Correctness of the customary assignment $x := t$ to normal variables x is captured by the ASSIGNMENT axiom that for a given postcondition P calculates the precondition by applying the substitution $[x := t]$ to P . In object-oriented programs one can additionally use assignments to instance variables and use the dereferencing $s.m(t)$ in the method calls. The latter calls for an extension of the assertion language with the corresponding expressions $e.u$.

F.S. de Boer proposed in [dB99] the following textually identical axiom for such assignments:

ASSIGNMENT TO INSTANCE VARIABLES

$$\{P[u := t]\} u := t \{P\},$$

where u is a (possibly subscripted) instance variable in *IVar*.

So, as in the case of the assignment to subscripted variables, the solution relies on an extension of the definition of substitution, in this case to the instance variables.

As usual, the definition of the substitution $P[u := t]$ proceeds by induction on the structure of P . The difference appears at the level of expressions s for which the definition of $s[u := t]$ is more elaborate. If s is of the form $e.u$ for an object expression e , the substitution has to take care of possible aliases of $e.u$. More precisely, if after applying the substitution inductively to e , the result refers to the currently active object, i.e., if $e[u := t] = \mathbf{self}$, then the outcome of the substitution is t . Otherwise, u is left untouched and the substitution is applied inductively to e . This is expressed in the following definition:

$$e.u[u := t] \equiv \mathbf{if} \ e[u := t] = \mathbf{self} \ \mathbf{then} \ t \ \mathbf{else} \ e[u := t].u \ \mathbf{fi}.$$

In case of only one (possibly recursive) method definition, the RECURSION rule for procedure calls can be adapted to method calls as follows, where we use the block statement discussed in Sect. 5.2:

RECURSION IV

$$\frac{\{P\} s.m(\mathbf{t}) \{Q\} \vdash \{P\} \mathbf{begin\ var\ self,\ u := s, \mathbf{t};\ S\ end\} \{Q\}}{\{P\} s.m(\mathbf{t}) \{Q\}}$$

where the method m is declared by $m(\mathbf{u}) :: S$.

So, as for recursive procedures we may use the desired conclusion as a hypothesis in the correctness proof of a block statement, where **self** and the list \mathbf{u} of formal parameters of the method are treated as local variables that are respectively initialised by the called object s and the list \mathbf{t} of actual parameters and are accessed in the method body S . A rule analogous to the RECURSION II rule is adopted to deal with total correctness.

To adjust correctness formulas that deal with generic method calls $y.m(\mathbf{x})$ to specific objects s and lists of actual parameters \mathbf{t} , we modify the SUBSTITUTION rule as follows, where we refer to the given set D of method definitions:

INSTANTIATION

$$\frac{\{P\} y.m(\mathbf{x}) \{Q\}}{\{P[y, \mathbf{x} := s, \mathbf{t}]\} s.m(\mathbf{t}) \{Q[y, \mathbf{x} := s, \mathbf{t}]\}}$$

where y, \mathbf{x} is a list of variables in Var which do not appear in D and $var(s, \mathbf{t}) \cap change(D) = \emptyset$.

In [AdBO09] these axioms and proof rules were used to establish total correctness of the above example program expressed by the correctness formula (4) and of an object-oriented program that inserts an element into a linked list.

9.3. Advanced topics in the verification of object-oriented programs

Various other aspects of object-oriented programming were studied from the correctness point of view. One of them is object creation. An object can be dynamically created by the assignment statement $u := \mathbf{new}$, where u is an object variable and **new** is a keyword in the considered programming language. The execution of this statement creates a *new* object and assigns its identity to the variable u . This new object comes with a default initialisation of all its instance variables. This can be modelled by some bookkeeping of the set of objects that are currently created, for example by maintaining a counter referring to an unbounded array. This allows one to reason about this assignment using a small program transformation. A drawback of this approach is that it refers to an explicit implementation.

An alternative is to use a substitution of $[x := \mathbf{new}]$ for object variables and define its application to a limited class of assertions (we call them *pure*) that take into account that object variables can only be compared for equality or be dereferenced in the method calls $s.m(\mathbf{t})$, and in which one does not quantify over such object variables. This leads to the following axiom proposed in [dB99]:

OBJECT CREATION

$$\{P[x := \mathbf{new}]\} x := \mathbf{new} \{P\},$$

where x is a simple (so not subscripted) object variable and P is a pure assertion.

We omit the details of the definition of the substitution $P[x := \mathbf{new}]$ and only remark that one cannot simply replace x in P by the keyword **new** because it is *not* an expression of the assertion language. The details are given in [AdBO09, Chapter 6].

C. Pierik and F.S. de Boer showed in [PdB03] how the approach of [dB99] can be extended to deal with classes, inheritance and subtyping. The same authors introduced in [dBP03] a general methodology for obtaining relatively complete Hoare's logics for object-oriented programs. A key issue is the extension of Gorelick's most general formulas [Gor75] (see Sect. 6.1) to deal with the states of object-oriented programs.

As another contribution to this line of research on verification of object-oriented programs let us mention [ÁdBdRS05] in which a Hoare-like proof system for partial correctness and deadlock freedom was developed for a subset of Java. The considered subset comprised the object-oriented core of Java, as well as concurrency via thread classes, allowing for a multithreaded flow of control. The Java concurrency model includes synchronous

message passing, dynamic thread creation, shared-variable concurrency via instance variables, and coordination via reentrant synchronization monitors. The verification method was formulated in terms of proof outlines that were tested both for interference freedom of shared-variable concurrency, as in [OG76a], and for cooperation of synchronous message passing, as in [AFdR80]. The authors established the soundness and relative completeness of their proof system. From an annotated program, a number of verification conditions were generated and discharged using the interactive theorem prover PVS [ORS92].

The verification of Java programs was also in the focus of tool-supported research projects around the year 2000. The LOOP project led by B. Jacobs considered sequential Java programs with specifications written in the Java Modeling Language (JML) [JP04]. JML was developed by a group led by G.T. Leavens [LCC⁺05]. It was inspired by the idea of *Design-by-Contract* introduced by B. Meyer in the context of the object-oriented programming language Eiffel [Mey97]. A contract is a specification in the form of assertions. The contract is agreed upon before an implementation is developed that should satisfy this contract. In its simplest form, a JML specification of a method comprises a precondition, one postcondition for normal termination and second one for exceptional termination, and a set of variables that may be changed. The LOOP compiler takes as input a Java program and a JML specification and outputs files in the syntax of the PVS theorem prover that describe the semantics of the program and proof obligations generated from the specification. These are then proven interactively using PVS.

For a subset of sequential Java programs, called NanoJava, a Hoare's logic was proven sound and relatively complete w.r.t. an operational semantics of the programs in the interactive theorem prover Isabelle/HOL [vON02]. In this approach, as in the paper [vO99] cited in Sect. 6.1, a semantic view of assertions was taken, whereby an assertion is a function from states to Boolean values. No separate syntactic representation of assertion was introduced. Thus, subtleties like the right definition of syntactic substitutions were not considered.

For the object-oriented programming language Eiffel a sound and relatively complete Hoare's logic was established in [NCMM09]. The emphasis was on the treatment of features in Eiffel that are not present in other object-oriented languages. These concern the details of the exception handling, so-called *once routines*, and multiple inheritance.

More recently B. Engelmann [EO16, Eng17] considered object-oriented programs in the context of *dynamic typing*. This means that variables do not have an a priori declared static type but may assume any value during the execution of the program. As an example consider the following method from [Eng17], where it is assumed that b is a Boolean variable:

$$\text{num_or_string}(b) :: \text{if } b \text{ then } x := y := 5 \text{ else } x := \text{'foo'}; y := \text{'bar'} \text{ fi}; z := x + y$$

If b is true, the method call $\text{num_or_string}(b)$ yields the numeric value $z = 10$, but if b is false, it yields the string value $z = \text{'foobar'}$. So the type of x and y is determined dynamically during runtime, and it is either numeric or string, with the operation $+$ being either addition or string concatenation. Such dynamically typed programs are present in the list processing language LISP and in widespread programming languages like Python and JavaScript.

Engelmann developed in [Eng17] a Hoare-like proof system for a model language **dyn** of object-oriented programs with dynamic typing. He proved soundness and relative completeness of his system, thereby extending the arguments of [Coo78, Gor75] and [dBP03] to the object-oriented part.

10. Alternative Approaches

10.1. Weakest precondition semantics and systematic program development

Dijkstra suggested an alternative approach to program verification, called *weakest precondition semantics* [Dij75]. The idea is that given a program S and a desired postcondition P we would like to find the weakest precondition $wp(S, P)$ such that $\{wp(S, P)\} S \{P\}$ holds in the sense of total correctness. 'Weakest' means here that for any precondition P such that $\{P\} S \{Q\}$ holds in the sense of total correctness, the implication $P \rightarrow wp(S, Q)$ holds.

This approach differs from Hoare's original approach by

- insisting on total correctness instead of on partial correctness,

- assuming that initially only the postcondition is given.

Additionally, Dijkstra insisted that the program should be developed *together* with its correctness proof. In [Dij75] he advanced this approach for his guarded command language that we briefly discussed in Sect. 7.1. Its notable feature was the use of programming constructs that support nondeterminism. Another interesting feature of the language was *parallel assignment* $\mathbf{x} := \mathbf{t}$, where \mathbf{x} is a list of different variables and \mathbf{t} is a list of expressions of the same length. This construct is for example useful to swap the values of variables without additional variables:

$$x, y := y, x.$$

The weakest precondition $wp(S, Q)$ is defined by induction on the structure of the program, with

- $wp(\mathbf{x} := \mathbf{t}, Q) \equiv Q[\mathbf{x} := \mathbf{t}]$,
- $wp(S_1; S_2, Q) \equiv wp(S_1, wp(S_2, Q))$,

as typical clauses.

The main problem is how to deal with loops. In [Dij75] the weakest precondition for guarded commands was defined as follows. Denote the alternative command $\mathbf{if} B_1 \rightarrow S_1 \square \dots \square B_n \rightarrow S_n \mathbf{fi}$ by IF , the repetitive command $\mathbf{do} B_1 \rightarrow S_1 \square \dots \square B_n \rightarrow S_n \mathbf{od}$ by DO and abbreviate $\bigvee_{i=1}^n B_i$ to BB . Then

$$wp(DO, Q) \equiv \exists k : k \geq 0 : H_k(Q),$$

where

$$H_0(Q) \equiv (Q \wedge \neg BB)$$

and for $k > 0$

$$H_k(Q) \equiv wp(IF, H_{k-1}(Q)) \vee H_0(Q).$$

Intuitively, $H_k(Q)$ is the weakest precondition guaranteeing proper termination in a state satisfying Q , after at most k guard selections.

Since k is used as a subscript of H_k , ' $\exists k : k \geq 0$ ' is here not a customary quantification but a shorthand for an infinite disjunction, with additionally the formulas $H_k(Q)$ defined by induction. In other words, so defined weakest precondition of a repetitive command is not an assertion in a first-order language. The same is the case for the weakest liberal preconditions mentioned in Sect. 4.3, that can be defined for **while** programs in an analogous way.

On the other hand, if one confines one's attention to the language of Peano arithmetic, then the weakest (liberal) preconditions can be encoded as first-order formulas, see for example [Win93, Chapter 7]. However, the resulting formulas are not natural and using them to express loop invariants (recall from Sect. 4.3 that weakest liberal preconditions can be used in a simple way to express loop invariants) is an overkill. This underlies the difficulty of finding simple loop invariants, the problem we already mentioned in Sect. 3.1, and clarifies why finding them is an important problem in developing correct programs or in establishing their correctness.

These complications disappear if one identifies assertions with the sets of states that satisfy them and views the weakest precondition as a way of assigning inductively semantics to programs. The above clauses defining wp show then how a program can be viewed as a *predicate transformer*, a concept further discussed in Sect. 10.3. So the weakest precondition approach can be viewed both as an approach to program verification and as a way of defining program semantics.

The concept of weakest precondition was used a number of times in this survey, which shows how this notion permeated research on Hoare's logic and program verification. The idea of using it to develop programs together with their correctness proofs was subsequently presented in book form in [Dij76a], where the weakest precondition semantics was extended to blocks and procedures. A systematic development of provably correct programs was further advanced by David Gries in his book [Gri81], notably by proposing a number of heuristics for finding loop invariants. Other books devoted to this subject are [Bac86, DF88, Kal90, BvW08]. Various aspects of the weakest precondition semantics were further discussed in book form in [DS90]. The problem of finding loop invariants became central in the subsequent study of program verification and development. [FMV14] surveyed various ways of constructing loop invariants, and provides their classification and analysis.

Another matter relevant for a systematic program development is termination. Both the WHILE II and RECURSION II rules provide only a kind of template for an actual termination proof, without explaining how

the termination functions are to be found. Some sophisticated techniques were developed to establish termination. They go beyond the framework of Hoare's logic or the weakest precondition semantics and rely on various methods developed in other areas, notably term rewriting systems. The authors of [CPR11] provided an accessible account of the recent developments. In particular, they explain a recent alternative approach to proving termination called *disjunctive termination argument*: only one of a disjunction of the termination functions needs to decrease, but this has to be the case after any number of iterations of the loop. It is argued that disjunctive termination arguments are easier to find than the classical termination argument dating back to Turing [Tur49] (see Sect. 2.1), where a single termination function has to decrease its values taken from a well-founded set with each iteration of the loop.

One difficulty of this approach for program verification is that it is not clear how to extend the weakest precondition semantics to deal with advanced programming constructs, for instance arbitrary procedures as considered in Sect. 6. In contrast, Hoare's approach is more flexible. In particular, as we have seen in Sect. 5.1, the reasoning can be supported by various adaptation rules. This cannot be done in the framework of the weakest precondition semantics which requires computing a single assertion. On the other hand, as we saw in Sect. 7.3, the weakest precondition semantics turned out to be helpful for the verification of probabilistic programs.

10.2. Specifying in Hoare's logic

Until now we used Hoare's logic exclusively to reason about programs. But the formalism of correctness formulas can be also used to specify programs, though some care needs to be exercised. For example, the correctness formula $\{\text{true}\} S \{z = x + y\}$ does not properly specify that a still to be constructed program S computes in z the sum of x and y because also $S \equiv x := 0; y := 0; z := 0$ satisfies this correctness formula.

To overcome this problem, one specifies the set $\text{var}(S)$ of variables that S may access. Then every variable outside $\text{var}(S)$ is known to keep its value during the computation of S . These variables can be used to freeze values of the initial state to use them in the postcondition evaluated in the final state. For example, by postulating $\text{var}(S) = \{x, y, z\}$, the correctness formula

$$\{x = x_0 \wedge y = y_0\} S \{x = x_0 \wedge y = y_0 \wedge z = x_0 + y_0\}$$

specifies that S computes in z the sum of x and y while leaving x and y unchanged. This specification technique is also at the heart of the most general formulas introduced by Gorelick in his completeness proof for recursive procedures [Gor75] and discussed in Sect. 6.1.

In a similar vein the adaptation rules keep track of the set $\text{var}(S)$ of the variables accessed by a program S , sometimes refined into the variables that are changed and those that are only read by S .

To support the specification of programs, variants of Hoare's logic have been proposed that introduce distinguished variables that may never be accessed by programs. By different authors, these variables are called logical variables (in [vON02]), logical constants (in [Mor94], further discussed in the next subsection), specification-only variables, or auxiliary variables (in [Kle99]). The last name, however, is in conflict with the concept of auxiliary variable introduced by Owicki and Gries in the context of parallel programs [OG76a] and discussed in Sect. 8.1.

T. Kleymann developed in [Kle99] a variant of Hoare's logic where such explicit logical variables appear in the assertions and a new consequence rule is proposed that takes care of these variables. As a result, he showed that this consequence rule is powerful enough to achieve adaptation completeness (a notion discussed in Sect. 5.1) without further adaptation rules.

We mentioned auxiliary variables first in the context of the Owicki-Gries method, when discussing verification of parallel programs in Sect. 8.1. The need for auxiliary variables in correctness proofs was already observed in [Cli73]. In [AdBO09] it was noted that the AUXILIARY VARIABLES rule is already needed to reason about disjoint parallel programs. Indeed, the correctness formula

$$\{x = y\} [x := x + 1 \parallel y := y + 1] \{x = y\}$$

cannot be proved using the DISJOINT PARALLELISM rule and the rules of the proof system \mathcal{H} .

The most extensive analysis of these matters was provided in [dGR16] that clarified and extended initial results of [Kle98] and [Kle99]. The authors of [dGR16] showed that the AUXILIARY VARIABLES rule is a *derived* rule both in the proof system \mathcal{H} and in a proof system for parameterless recursive procedures, which means that it can be eliminated from any proof that uses it. In the case of disjoint parallel programs the authors showed that this rule can be replaced by the simpler \exists -INTRODUCTION rule discussed in Sect. 5.1.

They also showed that for parallel programs with shared variables the AUXILIARY VARIABLES rule turns out to be essential. On the other hand, no other proof rules are needed. Indeed, Owicki proved in [Owi76] that the above presented proof system for parallel programs is complete in the sense of Cook.

10.3. Programming from specifications

R.-J. Back [Bac80] and C. Morgan [Mor94] extended Dijkstra's approach to a methodology for "programming from specifications" by a rule-based step-by-step development of specifications into programs. To this end, they introduced a *specification* statement. In [Mor94] it takes the form $\mathbf{x} : [P, Q]$, with the meaning

"If the initial state satisfies the precondition P then change only the variables listed in \mathbf{x} so that the resulting final state satisfies the postcondition Q ."

While the variables that are changed by a specification $S \equiv \mathbf{x} : [P, Q]$ are explicitly mentioned, namely $\text{change}(S) = \{\mathbf{x}\}$, there is no information which variables are accessed, i.e., $\text{var}(S)$ is not defined. Recall that various Hoare-style proof rules for programs S have application conditions concerning $\text{var}(S)$. Also, Gorelicks's most general formulas require fresh variables outside of $\text{var}(S)$ to freeze their initial values. To overcome this weakness, Morgan considers what he calls *logical constants* (occasionally also called *specification-only variables*), that by definition appear only in assertions and are thus never changed by any program or specification. With a logical constant X the initial values of a variable x can be frozen. For example, $x : [x = X, x > X]$ specifies that x should be increased. The amount of the increment is left unspecified.

The idea is that specifications and programs are handled on the same footing, so that programming operators like sequential composition or loops can be applied to specifications, as well. Constructs S_1 and S_2 in this extended syntax can be compared by a *refinement relation*: $S_1 \sqsubseteq S_2$ denotes that S_1 is refined by S_2 . Semantically, specifications and programs are considered as *predicate transformers* that transform given postconditions into the corresponding weakest preconditions. Refinement $S_1 \sqsubseteq S_2$ means that for all postconditions Q the implication

$$wp(S_1, Q) \rightarrow wp(S_2, Q)$$

holds, i.e., S_2 establishes the postcondition Q in at least all states where S_1 establishes Q . For example, $x : [x = X, x > X] \sqsubseteq x := x + 42$. In this approach specifications are stepwise *refined* to programs by the applications of refinement rules, described in [Mor94].

Morgan [Mor94] also considers a number of advanced programming concepts like recursive procedures with parameters, modules with local declarations of variables and procedures, and data refinement.

Subtly different from Morgan's specification statement is the *generic command* of J. Schwarz [Sch77]. It is written as $[P \Rightarrow Q][X]$, where P is a precondition, Q a postcondition, and X a set of variables. Semantically, a generic command denotes a certain state transformer. Assume an interpretation I . A *state transformer* T based on a finite set X of variables is a binary relation on the set Σ of states that has only read or write access to the states via the variables in the set X . For each program S its meaning is a state transformer based on $\text{var}(S)$.

Let $\mathcal{T}(X)$ be the set of all state transformers based on X . State transformers are ordered by the set inclusion: $T_1 \subseteq T_2$ means that T_1 produces less output than T_2 . The least element is the empty state transformer \emptyset producing no output, corresponding to a program that never terminates. Now, the semantics \mathcal{M}_I relative to I of a generic command $[P \Rightarrow Q][X]$ is given by

$$\mathcal{M}_I[[P \Rightarrow Q][X]] = \bigcup \{S \mid S \in \mathcal{T}(X) \text{ and } \{P\} S \{Q\} \text{ is true in } I \text{ in the sense of partial correctness}\}.$$

So given an interpretation I , $[P \Rightarrow Q][X]$ denotes the largest state transformer S based on X that satisfies $\{P\} S \{Q\}$ in the sense of partial correctness. 'Largest' refers here to the ordering of state transformers by the set inclusion. In [Old83b] it was shown that for $\{\mathbf{x}\} = X$ the formula W of the ADAPTATION II rule, introduced in Sect. 5.1, expresses the weakest liberal precondition $wlp_I([P \Rightarrow Q][X], R)$ of the generic command $[P \Rightarrow Q][X]$ w.r.t. a given postcondition R . This connects generic commands via the weakest liberal preconditions to the ADAPTATION II rule.

Schwarz used generic commands for stating an alternative version of a RECURSION rule, but did not embark on program development. Morgan's specification statement $\mathbf{x} : [P, Q]$ is similar to the generic command $[P \Rightarrow Q][X]$, except that the list \mathbf{x} does not record the variables that are only read, which are also covered by X .

10.4. Algorithmic logic and dynamic logic

Hoare's logic is geared towards establishing program correctness. However, from the point of view of mathematical logic it has a very rigid syntax: the correctness formulas cannot be negated or combined, for example by disjunction. As a result one cannot view Hoare's logic as an extension or a modification of some existing logics, even though it crucially relies on first-order logic and its extensions, for example by allowing subscripted variables. In some alternative approaches one could view reasoning about programs as an extension of reasoning within existing logics. We discuss now briefly two most prominent examples.

Algorithmic logic was originally proposed in [Sal70] and presented in book form in [MS87]. It extends first-order language by expressions that can be interpreted as programs and constructs that allow one to mix formulas and programs. Interestingly, substitutions for a sequence of variables are viewed as atomic programs. This is equivalent to the aforementioned parallel assignment of Dijkstra, and also shows a close connection with the ASSIGNMENT axiom. Programs are built by allowing formulas as tests and using program composition, conditionals, and loops, all written in a compact notation. For example, **if** B **then** S **else** T is written as $\underline{v}[BST]$. In turn, the construct $*[BST]$ corresponds to **while** B **do** S **od**; T .

Expressions of the form $S\phi$, where S is a program and ϕ a formula, correspond to the strongest postcondition introduced in Sect. 4.3. Further, it is shown how termination of the considered programs can be expressed as a formula that admits countable disjunction. This is analogous to the definition of $wp(DO, Q)$ given earlier.

The relation between Hoare's logic and algorithmic logic becomes clear when one realizes that the correctness formula $\{P\} S \{Q\}$ can be expressed as the implication $SP \rightarrow Q$. Consequently, rules used in Hoare's logic can be readily reproduced as rules in algorithmic logic, in particular various forms of adaptation rules. This straightforward modelling, however, does not yield new insights concerning program verification.

Research on algorithmic logic focused mostly on such matters as studies of consistency and the infinitary completeness of selected theories, derivation of the normal forms of programs, and axiomatization of various data structures, rather than on (relatively) complete axiomatizations of fragments concerned with specific features of programming languages, a direction Hoare's logic took.

Dynamic logic was originally proposed in [Pra76] and presented in book form in [Har79] and more extensively in [HKT00]. It is very similar to algorithmic logic introduced six years earlier, though it was developed independently. It enriches first-order logic by constructs reminiscent of modal logic. First, programs are defined starting from atomic actions and tests, using the sequential composition ($;$), nondeterministic composition (\cup) (absent in algorithmic logic), and iteration ($*$) that corresponds to Kleene's star. In the dynamic logic syntax the **while** B **do** S **od** statement can be expressed as $(B; S)^* \cup \neg B$.

Further, one admits formulas of the form $[S]\phi$, where S is a program and ϕ is a formula, with the intended interpretation "every execution of the program S from the current state leads to a state in which ϕ is true". The formulas and programs are defined by simultaneous induction, allowing the usual Boolean connectives. A dual formula to $[S]\phi$ is $\langle S \rangle \phi$, defined by:

$$\langle S \rangle \phi \equiv \neg[S]\neg\phi.$$

So its intended interpretation is "some execution of the program S from the current state leads to a state in which ϕ is true". The $[S]$ and $\langle S \rangle$ operators can thus be viewed as the counterparts of the \square and \diamond operators in propositional modal logic, but parameterized with a program S .

Research on dynamic logic mainly focused on a study of various fragments or extensions, with the corresponding sound and complete axiomatizations, and the corresponding decidability and computational complexity results. In particular, dynamic logic was extended in [Har79] to deal with recursive procedures.

Typical axioms are:

$$[S; T]\phi \leftrightarrow [S][T]\phi,$$

which corresponds to one of the mentioned clauses that define the weakest precondition, and

$$[S^*]\phi \leftrightarrow \phi \wedge [S][S^*]\phi,$$

that captures the idea that $*$ stands for the infinite iteration.

The relation between Hoare's logic and dynamic logic is easily established by noticing that the correctness formula $\{P\} S \{Q\}$ can be expressed as the implication $P \rightarrow [S]Q$. So $[S]Q$ models what is called the weakest liberal precondition introduced in Sect. 4.3. Consequently, as in the case of algorithmic logic, proof rules of the proof system \mathcal{H} of Sect. 3.1 can be translated into proof rules of dynamic logic. Moreover, these translated rules are derivable from the adopted axioms and proof rules of dynamic logic.

Further, thanks to the richer syntax, it is possible to express other program properties and discuss such properties like program equivalence. For example, the following formula states that the program S is deterministic:

$$\langle S \rangle \text{true} \rightarrow [S] \text{true}.$$

Dynamic logic was generalized in a number of ways, for example to *epistemic dynamic logic* (see, e.g., [BR16]), to allow reasoning about modalities and change. Also it was extended to specify and reason about *hybrid systems*, i.e., systems where a discrete control interacts with a continuous dynamics [Pla08]. Further, the KeY system (discussed Sect. 11) is based on dynamic logic.

10.5. Temporal logic and model checking

Research on verification of parallel programs carried out in the seventies showed limitations of Hoare's logic in reasoning about concurrent programs. We discussed three properties of such programs: partial correctness, termination, and absence of deadlock. However, in contrast to the sequential programs, concurrent programs are often supposed to operate repeatedly, in a cyclic fashion. For instance, a solution to the mutual exclusion problem deals with infinite executions of the program components operating in parallel. This calls for a study of properties (such as an *eventual access*) that cannot be expressed in Hoare's logic.

To express such concepts and to systematically reason about them A. Pnueli proposed in [Pnu77] to use *temporal logic*. Using it one can express in a natural way various program properties that do not necessarily deal with the input/output behaviour of a program. In contrast to Hoare's logic the reasoning about programs is not syntax-directed. Instead, one usually reasons about specific control points in a program and a relation between them.

For example, the following formula states that a process P infinitely often enters its critical section CS :

$$\Box \Diamond \text{ in } CS$$

where *in CS* is a formula that states that the control in the process P is within CS .

In turn, to express the strong fairness assumption for the repetitive command

$$S \equiv \mathbf{do} B_1 \rightarrow S_1 \Box \dots \Box B_n \rightarrow S_n \mathbf{od}$$

we can use the following formula:

$$\bigwedge_{i=1}^n (\Box \Diamond (at S \wedge B_i) \rightarrow \Box \Diamond at S_i),$$

where *at T* holds when the control in the considered program is just in front of T .

Appropriate axioms and proof rules were then developed to reason about such formulas. Temporal logic, applied to concurrent programs, grew into an impressive research area, see in particular the books [MP91, MP95] of Z. Manna and A. Pnueli. This line of research should be viewed as complementary to Hoare's logic, that is why we do not devote more space to it.

In that context a distinction, first advanced in [Lam77], is useful. A *safety property* states 'nothing bad will happen' during a program execution, while a *liveness property* states that eventually 'something good will happen' during a program execution. According to this distinction partial correctness, absence of errors, and deadlock freedom are safety properties. In contrast, program termination, fairness, eventual access, infinite access, etc. are liveness properties. In temporal logic these properties can be formulated by means of invariants that become formulas of the form $\Box \phi$, while liveness properties are formulas of the form $\Diamond \phi$. Hoare's logic is a convenient vehicle to prove safety properties. Its modification to deal with termination is also quite natural. As we saw in Sect. 7.2 one can also investigate within this logic fairness of nondeterministic programs. In fact, even more advanced liveness properties, such as fairness of parallel programs and eventual access properties can be studied,

as well (see [OA88]). However, the treatment becomes awkward, as one has to reason then about appropriately transformed programs. This is in contrast to temporal logic that allows one to reason about liveness properties in a simple way, without any need for program transformations.

Temporal logic led in turn to *model checking*. Independently, E.M. Clarke and A. Emerson [EC82], as well as J.P. Queille and J. Sifakis [QS81], discovered that the problem of checking whether a finite-state system satisfies (is a *model* of) a propositional temporal logic formula is decidable, providing efficient algorithms for it. This was the start of enormous research activities extending the scope of model checking so that even industrial-size problems could be tackled [CGH⁺93]. Interestingly, model checking is often used to debug a system because in case the system does not satisfy the temporal logic specification, model checkers can provide a counterexample that is helpful for understanding the mismatch between a system and its specification. In the recent years model checking has been extended to infinite-state systems, mostly by automatically constructing and refining abstractions of the system, a method known as *counterexample-guided abstraction refinement*, see [CGJ⁺03]. The state of the art of model checking is represented in the handbook [CHVB18].

Another approach based on temporal logic is *Temporal Logic of Actions* (TLA) introduced by L. Lamport in [Lam94]. It is a method of specifying and reasoning about concurrent systems in which systems and their properties are represented in the same temporal logic. A system is specified by a temporal formula representing a single loop that in each iteration nondeterministically chooses one action for execution. An action describes a basic step of the considered system and is specified by a predicate using normal and primed variables as we have seen it already in Turing's example in Sect. 2.1. The standard form of a TLA specification is

$$Init \wedge \Box [Next]_x \wedge Live,$$

where *Init* is the initial predicate, *Next* is the next-state relation expressed as a disjunction of actions, *x* is the list of all variables occurring in these actions, and *Live* is a temporal formula that specifies a liveness condition in terms of fairness assumptions. To allow for refinement of concurrent systems in the presence of different granularities of the basic steps, *stuttering steps* are considered. To hide parts of the state, quantification over variables is possible. The extension TLA+ comes with facilities to structure system specifications into modules, see [Lam02]. Properties of a subset of TLA+ specifications can be verified with the model checker TLC discussed in [Lam02].

10.6. Separation logic

This approach to program verification was originally developed as an extension of Hoare's logic to reason locally about pointer structures, see [ORY01, Rey02, OYR04, O'H19]. To cope with pointers, separation logic builds upon a semantic model, in which a state is a pair (s, h) consisting of a *store* s and a *heap* h . A store s is a mapping from variables to values (so a state in the sense of Sect. 4.2), which may be data values such as integers, or pointer values such as addresses. A heap h is a finite partial mapping from addresses (or cells) to values, which again can be data or pointers. It is assumed that the addresses are integers, which in turn are values.

Several low-level statements for manipulating the heap were considered in [ORY01, O'H19]. Let x stand for a variable and e for an integer expression, which can denote an address in the heap. One can explicitly distinguish between an address and its contents: $[e]$ denotes the contents of the heap at address e . *Assignments* affect only the store. Besides the normal assignments $x := e$ there are *lookups* $x := [e]$, where e is interpreted as an address in the heap and the contents of e is assigned to the variable x in the store. *Mutations* affect the heap. An *update* $[e] := e'$ expresses that the contents of the address e in the heap becomes value of the expression e' . Further, $x := \text{alloc}()$ expresses that the address x is newly allocated to the heap, and $\text{free}(x)$ expresses that the address x is deallocated from the heap.

Separation logic extends the usual assertion language of Hoare's logic to specify properties of the heap: emp asserts that the heap is empty, $e \mapsto e'$ asserts that the heap consists of one cell, with address e and contents e' , and $e \mapsto -$ asserts that the heap consists of one cell, with address e but unknown content. The main new operator in separation logic is the *separation conjunction*, written $*$ and pronounced "and separately". The assertion $P * Q$ expresses that the heap can be split into two disjoint parts in which P and Q hold, respectively. Using separation conjunction, one can specify larger parts of the heap. For instance, the assertion $(x \mapsto 21) * (y \mapsto 42)$ describes two separate cells with addresses x and y and contents 21 and 42, respectively. As abbreviation one uses $e \mapsto e_1, \dots, e_n$ to stand for $e \mapsto e_1 * e + 1 \mapsto e_2 * \dots * e + n - 1 \mapsto e_n$ and thus asserting that the heap consists of n adjacent cells with addresses $e, \dots, e + n - 1$ and contents e_1, \dots, e_n , respectively. For example, the

assertion $(x \mapsto 21, y) * (y \mapsto 42, x)$ concisely specifies a heap with a cyclic pointer structure, in which $x \mapsto 21$, y stands for $x \mapsto 21 * x + 1 \mapsto y$ and similarly with $y \mapsto 42$, x . It consists of two separate parts of the heap at the addresses x and y that contain 21 and 42, respectively, and a pointer to the other address.

For mutations and lookups the following axioms were stated in [Bro07, O'H19]:

ALLOCATION

$$\{\mathbf{emp}\} x := \mathbf{alloc}() \{x \mapsto -\}$$

DE-ALLOCATION

$$\{x \mapsto -\} \mathbf{free}(x) \{\mathbf{emp}\}$$

UPDATE

$$\{e \mapsto -\} [e] := e' \{e \mapsto e'\}$$

LOOKUP

$$\{P[x := e'] \wedge e \mapsto e'\} x := [e] \{P \wedge e \mapsto e'\}$$

where x does not occur in e or e' .

Separation conjunction enables the formulation of proof rules for local reasoning about components of parallel programs. Crucial is the following rule, which is essentially the INVARIANCE rule in a semantic setting of heaps:

FRAME

$$\frac{\{P\} S \{Q\}}{\{P * R\} S \{Q * R\}}$$

where $\mathit{free}(R) \cap \mathit{change}(S) = \emptyset$.

In separation logic, this rule serves to extend a local specification involving only the variables and parts of the heap that is used by S by adding assertions about variables and other parts of the heap not modified by S . Thus the FRAME rule is considered as the key to local reasoning about the heap [Rey02].

Separation logic was originally used for the verification of sequential programs manipulating pointer structures. Exploiting its ability to reason explicitly about the heap, the approach was later extended in [O'H07] to reason in a modular way about concurrent programs. The idea is that two threads that operate on disjoint parts of the heap do not interfere, and thus can be verified in isolation. This is captured by the following rule, which is the DISJOINT PARALELISM rule in a semantic setting with heaps:

CONCURRENCY

$$\frac{\{P_1\} S_1 \{Q_1\}, \{P_2\} S_2 \{Q_2\}}{\{P_1 * P_2\} [S_1 \parallel S_2] \{Q_1 * Q_2\}}$$

where $\mathit{free}(P_1, Q_1) \cap \mathit{change}(S_2) = \mathit{free}(P_2, Q_2) \cap \mathit{change}(S_1) = \emptyset$.

In this rule, the separation conjunction $P_1 * P_2$ in the precondition of the parallel composition $[S_1 \parallel S_2]$ is true if the heap can be partitioned into sub-heaps making the local preconditions of the components S_1 and S_2 true. The premises of the rule state that the components S_1 and S_2 establish local postconditions Q_1 and Q_2 , respectively, which in the conclusion are combined into the global postcondition $Q_1 * Q_2$.

A major issue was developing a coherent semantical model for this *concurrent separation logic*. This was solved by S. Brookes in [Bro07]. His semantics evaluates resource-sensitive partial correctness formulas of the form $\Gamma \vdash \{P\} S \{Q\}$, where Γ is a *resource context* that specifies for each resource name occurring in the program S a finite set of variables, a protection list, and a resource invariant, and a proof system allows one to reason about these formulas.

10.7. Relational Hoare logic

We conclude this overview of alternative approaches by discussing a line of research that began with the work of N. Benton, [Ben04]. In it Hoare's logic was modified to verify correctness of various optimizing program transformations. This was achieved by proposing a proof system in which one reasons at the same time about a pair of programs the variables of which are related by some precondition and postcondition. To this end the customary correctness formulas were replaced by *judgments* concerning two **while** programs, S_1 and S_2 . These are statements of the form

$$S_1 \sim S_2 : \Psi \Rightarrow \Phi,$$

where Ψ and Φ are relations on program states. Informally, such a judgment states that if initially the relation Ψ between the variables of the programs S_1 and S_2 holds, then after their independent executions the relation Φ between their variables holds. (So a notation $\{\Psi\} S_1 \sim S_2 \{\Phi\}$ would have been more intuitive.) Typically both programs use the same variables, so to indicate from which program the variable is taken the indices $\langle 1 \rangle$ and $\langle 2 \rangle$ are used in Ψ and Φ .

As an example consider the technique of *invariant hoisting*. The program

$$S_1 \equiv \mathbf{while} \ i < n \ \mathbf{do} \ x := y + 1; \ i := i + x \ \mathbf{od}$$

can be optimized to

$$S_2 \equiv x := y + 1; \ \mathbf{while} \ i < n \ \mathbf{do} \ i := i + x \ \mathbf{od}.$$

The resulting programs compute the same values for the variables i and n . This can be expressed as the judgment

$$S_1 \sim S_2 : \Phi \Rightarrow \Phi,$$

with $\Phi \equiv i\langle 1 \rangle = i\langle 2 \rangle \wedge n\langle 1 \rangle = n\langle 2 \rangle \wedge y\langle 1 \rangle = y\langle 2 \rangle$. So the first occurrence of Φ expresses the information that prior to the programs executions the values of their variables i, n, y are respectively equal, while the second occurrence states that the values of these variables are respectively equal after the executions of both programs.

In general, if we rename S_1 to S'_1 by appending to each of its variables the index $\langle 1 \rangle$, and similarly with S_2 for which we use the index $\langle 2 \rangle$, then the judgment $S_1 \sim S_2 : \Psi \Rightarrow \Phi$ can be interpreted as the correctness formula $\{\Psi\} S'_1; S'_2 \{\Phi\}$.

The resulting logic was called *Relational Hoare Logic* (RHL), where the qualification 'Relational' referred to the fact that instead of assertions relations were used. An example rule is the following one dealing with the conditional statement:

$$\frac{S_1 \sim S_2 : \Psi \Rightarrow \Phi}{\mathbf{if} \ B \ \mathbf{then} \ S_1 \ \mathbf{else} \ S_2 \ \mathbf{fi} \ \sim S_1 : \Psi \Rightarrow \Phi}$$

The premise of this rule states that in the context determined by the Ψ and Φ relations, the statement S_1 can be replaced by S_2 . Thus the second branch of the conditional statement can be replaced by S_1 , and consequently this conditional statement can be simplified to S_1 .

Appropriate proof rules formulated in this framework made it possible to verify various known compiler optimization techniques, including the above example of invariant hoisting.

In [BGB09] this approach was extended to study equivalence of probabilistic programs. The considered programs extend **while** programs by allowing a standard feature of probabilistic programs, called *random sampling*, which is an assignment of the form $x := \mathcal{D}$, where \mathcal{D} is a probability distribution over the values of the type of x . The semantics of such programs, as in [dHdV02], is then a mapping from the set of states to the set of distributions over the states.

The semantics of the judgments is defined in such a way that it respects the fact that in both considered programs the same probability distributions are used. In particular the judgment $S \sim S : \mathbf{true} \Rightarrow x\langle 1 \rangle = x\langle 2 \rangle$, where $S \equiv x := \mathcal{D}$, is true.

The resulting logic was called *probabilistic Relational Hoare Logic* (pRHL). pRHL is not a simple extension of RHL because in presence of probabilities some rules of RHL become unsound. For example, the rule

$$\frac{S_1 \sim S_2 : \Psi \Rightarrow \Phi_1, \ S_1 \sim S_2 : \Psi \Rightarrow \Phi_2}{S_1 \sim S_2 : \Psi \Rightarrow \Phi_1 \wedge \Phi_2}$$

is sound in RHL (because the programs are deterministic) but not in pRHL.

This work was carried out in the context of *CertiCrypt*, a system built on top of the interactive theorem prover Coq (see Sect. 11), that makes it possible to provide machine-checked correctness proofs of cryptographic algorithms.

In [BKOB13] this framework was further generalized to allow for a probabilistic reasoning about *differential privacy*, a notion of privacy that guarantees that the behaviour of an algorithm taking values from a database hardly changes when the database is slightly modified. Differential privacy in particular provides a formal guarantee that information about specific participants in a database is not revealed by the algorithm. The resulting formalism is called *approximate probabilistic Relational Hoare Logic* (apRHL).

In apRHL, the judgments of [BGB09] were generalized to parameterized judgments concerning two probabilistic programs. The appropriate proof rules generalize those of [BGB09]. This framework was implemented in a system called *CertiPriv* built on top of Coq, which was in particular used to provide machine-checked proofs of soundness of the considered rules. This work was further pursued in [BGA⁺14], where differential privacy was dealt with by means of a transformation of a probabilistic program into a non-probabilistic one that simulates two executions of the original program. As a result, differential privacy of a single program could be established using the original proof system \mathcal{H} presented in Sect. 3.1.

11. Final Remarks: a Summary and an Assessment

Hoare's logic had a huge impact on program verification, notably by allowing one to approach it in a systematic way, using the logical apparatus of formal proofs. Combining it with the research on program semantics made it possible to argue about the soundness and relative completeness of the underlying proof systems. The syntax-directed form of Hoare's logic suggested a natural research agenda, which —as we have seen— allowed one to deal with several programming constructs and forms of program construction, including higher-order procedures, nondeterminism, concurrency, and object-orientation.

This survey aimed at providing a systematic exposition of these developments. Because of space considerations we had to omit an account of various Hoare-style proof rules for such concepts as program jumps ([CH72]), go-to statement ([dB80, Chapter 10], written by A. de Bruijn), or several forms of abrupt loop termination present in Java ([HJ00]).

In the seventies Hoare's logic was used to define programming languages. In [Hoa72a] and [HW73] an axiomatic definition of the programming language PASCAL was given. These papers provided proof rules for simple constructs such as the **case** statement and the **repeat** statement. However, the presentation was incomplete. For example, no account of reasoning about recursive procedures or pointers was given. This work was pursued in [LGH⁺78] where axioms and proof rules in Hoare's logic for the programming language EUCLID were presented.

From the current perspective one can see that such axiomatic presentations were not rigorous since no soundness proofs were provided to justify the introduced axioms and rules, notably the recursion rule. To see that such soundness proofs are not superfluous recall from Sect. 5.1 the observation of [Old83b] that the adaptation rule for EUCLID is not sound. Also, termination was not dealt with in these papers and the reasoning about it can be another source of possible, subtle, errors (see for example the discussion at the end of Sect. 6.1). In fact, as we saw, reasoning about soundness within the framework of Hoare's logic started only after these two papers were published.

But even if such soundness proofs were presented, given the size of the considered programming languages, there would be a non-trivial chance of errors. In fact, we mentioned a number of times that arguments about correctness of various proof rules in Hoare's logic or about soundness or relative completeness of some specific proof systems have led to various, occasionally, pretty subtle errors.

A logical remedy is to use automated reasoning to argue about various Hoare's logics. This brings us to *computer aided verification*, one of the important research directions that emerged in the eighties. Since 1990 it is a subject of an annual conference with the same name (abbreviated to CAV). In relation to Hoare's logic this work naturally divides into two categories:

- mechanical verification by means of interactive theorem provers of properties of various Hoare logics, such as soundness and relative completeness proofs,
- computer aided verification of selected programs.

In what follows we refer to a number of systems that are the outcome of several years of research and associated ongoing implementations. Given that mechanical verification is not the subject of this survey, we do not discuss these systems in detail. In each system, to appropriately formalize the assertion-based reasoning about programs and reason about the corresponding proof systems, several important design decisions have to be taken concerning the choice of the underlying assertion language, proof representation, choice of proof tactics, form of the interaction with the user, etc. A separate survey on this subject would be a most welcome addition to this one.

An early contribution in the first category was [Sok87], where soundness of the original proof system of Hoare from [Hoa69] was established in LCF (that stands for Logic for Computable Functions), an early interactive automated theorem prover developed in the seventies, see, e.g. [GMW79].

Next, in [Kle98], [Nip02a] and [Nip02b] soundness and relative completeness of Hoare’s logics for partial and total correctness of **while** programs and programs with recursive procedures was established in the LEGO and Isabelle/HOL interactive theorem provers, the latter system described initially in [NPW02]. In Sect. 6.1, we mentioned already that in [vO99] soundness and relative completeness of a proof system for a programming language with the mutually recursive procedures with the call-by-value mechanism was established in Isabelle. In all these works assertions are identified with sets of states that satisfy them.

Further, in [NN99] the semantics and proof system of [OG76a] for partial correctness of parallel programs, that we discussed in Sect. 8.1, was formalized in the Isabelle/HoL system. Subsequently the authors proved soundness of this proof system and verified a number of correctness proof examples. An analogous formalization was carried out in [Bal06] on the basis of dynamic logic in the KIV system, described in [BRS⁺00]. This approach combined symbolic execution of the operational semantics of the programs with induction.

The works in the second category concern specific programs. They rely on various systems and tools that support mechanical verification of programs based on Hoare’s logic or some of the alternative approaches discussed in the previous section. An early example is [Gor88], where a simple program verifier implemented in LISP is described. It consists of a generator of verification conditions and a theorem prover. It was used in particular to justify Hoare’s correctness proof given in Fig. 4 in Sect. 3.1. The following more recent selective contributions underscore the importance of mechanical verification. Several other recent systems and tools for program verification are described in [PR18].

- In his paper [Hoa71b] concerned with the correctness of the FIND program, Hoare expressed the desire for computer support in “formulating the lemmas, and perhaps even checking the proofs.” Only much later, Filliâtre [Fil07] published a mechanized proof of FIND using the interactive theorem prover Coq and following Hoare’s proof from [Hoa71b] as closely as possible. He noticed that Hoare’s informal termination proof does not meet the requirements of a termination function in the sense that the additional invariants used by Hoare are not real invariants.

Coq is based on a formal language called the *Calculus of Inductive Constructions* that extends typed lambda calculus, see, e.g., [Ch13]. It was used to verify formally proofs of a number of famous mathematical results, including the Four Color theorem.

- In [CDE⁺16] a tool-supported correctness proofs of Quicksort and its variants were reported. The authors used the Hoare’s logic based verification tool Dafny [Lei10]. Dafny comprises (1) an imperative, class-based language for programs annotated with the features needed for verification (such as assertions, *framing* clauses expressing the objects that are allowed to be changed, *ghost variables*, i.e., auxiliary variables in the sense of Sect. 8.1, and termination functions) and (2) a verifier that translates annotated Dafny programs to an intermediate language input by the tool Boogie [BCD⁺05]. The latter is a fully-automatic deductive verification system that generates verification conditions which are then passed to an SMT (satisfiability modulo theories) solver, where the default SMT solver is Z3 [dMB08]. Boogie uses the weakest precondition semantics discussed in Sect. 10.1.
- TimSort is a standard sorting algorithm provided by several programming framework, including the Java standard library. While trying to verify it the authors of [dGdBB⁺19] discovered a bug. They formally specified a revised version and mechanically verified its correctness, including termination, in the KeY system. KeY is an extensive software development system that supports in particular specification and formal verification of object-oriented software, see [ABB⁺16]. As mentioned earlier, it is based on dynamic logic. The KeY-Hoare tool (see [BH16]), built on top of the KeY system, allows one to reason about partial and total correctness of **while** programs in an extension of Hoare’s logic with explicit state updates.

In recent years research on Hoare's logic visibly slowed down, probably due to the fact that through hundreds of publications it achieved its main goal of creating a comprehensive formal framework to reason about various classes of programs. On the other hand, its versatility results in new applications in various contexts. We provide here three recent examples.

The first, due to [AMO13], concerns reasoning about linear systems (for example linear differential equations) that are expressed graphically as block diagrams. The authors, building upon the framework of abstract Hoare's logic developed in [AMMO09] mentioned in Sect. 4.3, provided a sound Hoare-like proof system that allowed them to reason about various examples of linear systems, including linear filters and the steam boiler problem, a classic example of a hybrid system.

The second one concerns reasoning about time complexity of the programs. In [HN18] three Hoare's logics for reasoning about time bounds, including the original logic due to [Nie87], were formalized and shown to be sound and relatively complete.

The final one is an application to quantum programs, a research direction that originated with the work of M. Ying [Yin11]. For a recent overview of the developments on this subject, that successfully parallel the developments of the customary Hoare's logic, see [Yin19]. A related approach of [Unr19] follows the line of research started with [Ben04] and [BGB09] and introduces *quantum Relational Hoare Logic* (qRHL) that allows one to reason about how the outputs of two quantum programs relate given a relation between their inputs.

One should also mention here recent work aimed at supporting teaching of Hoare's logic. In [SS14] an account was given of a tool called HABA (Hoare Advanced Homework Assistant) that was specifically designed to teach Hoare's logic. The tool supports reasoning about **while** programs with integer variables and arrays.

Hoare's logic is an example of an *assertion-based* method of program verification, in the sense that it relies on the use of assertions. This is also a common feature of the approaches to program verification discussed in the previous section. In this sense we can say that all these approaches are direct or less direct successors of Turing's approach from [Tur49].

However, there are other theories of program verification, notably of parallel and distributed programs, that take a different approach. A prominent example is CCS (Calculus for Communicating Systems) of Robin Milner that focuses on reasoning about equivalence of abstract distributed programs [Mil80, Mil89]. It is at the origin of a number of related theories and led to the field of *process algebra* [BPS01]. A comparative survey of these approaches would form a natural complement to this one.

Acknowledgements

We would like to thank all three reviewers for exceptionally detailed and helpful referee reports that led us to improve and expand the presentation.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

A. Turing's example

We present here in Fig. 6 a proof outline for a **while** program corresponding to Turing's example given in Fig. 3 of Sect. 2.1. The program without the assertions was shown in [OW12]. The qualification **inv** is used to annotate the loop invariants and **bd** to annotate the termination functions. So this is a proof outline establishing total correctness of the program w.r.t. the precondition $n \geq 1$ and postcondition $v = n!$. The assertions used for partial correctness are similar to the ones used in the correctness proof of [dB75] of a program corresponding to Turing's example that uses both a **while** and a **repeat** statement.

```

{ $n \geq 1$ }
 $r := 1;$ 
 $u := 1;$ 
 $v := u;$ 
{inv :  $P_1 \equiv v = r! \wedge u = r! \wedge 1 \leq r \leq n$ }
{bd :  $t_1 \equiv n - r$ }
while  $r < n$  do
{ $P_1 \wedge r < n$ }
{ $v = r! \wedge u + v = 2 \cdot v \wedge 1 \leq 2 \leq r + 1 \leq n$ }
 $s := 1;$ 
 $u := u + v;$ 
 $s := s + 1;$ 
{inv :  $P_2 \equiv v = r! \wedge u = s \cdot v \wedge 1 \leq s \leq r + 1 \leq n$ }
{bd :  $t_2 \equiv r + 1 - s$ }
while  $s \leq r$  do
{ $P_2 \wedge s \leq r$ }
{ $v = r! \wedge u + v = (s + 1) \cdot v \wedge 1 \leq s + 1 \leq r + 1 \leq n$ }
 $u := u + v;$ 
 $s := s + 1$ 
{ $P_2$ }
od
{ $P_2 \wedge s > r$ }
{ $v = r! \wedge u = s \cdot v \wedge 1 \leq s = r + 1 \leq n$ }
{ $u = (r + 1) \cdot r! \wedge 1 \leq r + 1 \leq n$ }
{ $u = (r + 1)! \wedge 1 \leq r + 1 \leq n$ }
 $r := r + 1;$ 
 $v := u$ 
{ $P_1$ }
od
{ $P_1 \wedge r \geq n$ }
{ $v = r! \wedge r \leq n \wedge r \geq n$ }
{ $v = n!$ }

```

Fig. 6. Turing's example as a proof outline of a corresponding **while** program

References

- [ABB⁺16] Ahrendt W, Beckert B, Bubel R, Hähnle R, Schmitt PH, Ulbrich M (eds) (2016) Deductive software verification—the KeY book—from theory to practice, volume 10001 of Lecture notes in computer science. Springer
- [ABC87] Apt KR, Bougé L, Clermont P (1987) Two normal form theorems for CSP programs. *Inf Process Lett* 26:165–171
- [AdB77] Apt KR, de Bakker JW (1977) Semantics and proof theory of pascal procedures. In: Salomaa A, Steinby M (eds) Automata, languages and programming: proceedings of the fourth colloquium, volume 52 of Lecture notes in computer science. Springer, pp 30–44
- [AdB90a] America P, de Boer FS (1990) Proving total correctness of recursive procedures. *Inf Comput* 84(2):129–162
- [AdB90b] America P, de Boer FS (1990) A proof system for process creation. In: Broy M (ed) Programming concepts and methods: proceedings of the IFIP working group 2.2, 2.3 working conference on programming concepts and methods. North-Holland, pp 303–332
- [AdB19] Apt KR, de Boer FS (2019) Reasoning about call-by-value: a missing result in the history of Hoare's logic, [arXiv:1909.06215](https://arxiv.org/abs/1909.06215)
- [ÁdBdRS05] Ábrahám E, de Boer FS, de Roever WP, Steffen M (2005) An assertion-based proof system for multithreaded java. *Theor Comput Sci* 331(2–3):251–290
- [AdBO90] Apt KR, de Boer FS, Olderog E-R (1990) Proving termination of parallel programs. In: Feijen WHJ, van Gasteren AJM, Gries D, Misra J (eds) Beauty is our business, a birthday salute to Edsger W. Dijkstra. Springer, New York, pp 0–6
- [AdBO09] Apt KR, de Boer FS, Olderog E-R (2009) Verification of sequential and concurrent programs, 3rd edn. Springer, New York
- [AdBOdG12] Apt KR, de Boer FS, Olderog E-R, de Gouw S (2012) Verification of object-oriented programs: a transformational approach. *J Comput Syst Sci* 78(3):823–852
- [AFdR80] Apt KR, Francez N, de Roever WP (1980) A proof system for communicating sequential processes. *ACM Trans Program Lang Syst* 2(3):359–385
- [AFK88] Apt KR, Francez N, Katz S (1988) Appraising fairness in distributed languages. *Distrib Comput* 2(4):226–241
- [AFPdS11] Almeida JB, Frade MJ, Pinto JS, de Sousa SM (2011) Rigorous software development—an introduction to program verification. Undergraduate topics in computer science. Springer
- [AMMO09] Arthan R, Martin U, Mathiesen EA, Oliva P (2009) A general framework for sound and complete Floyd-Hoare logics. *ACM Trans Comput Log* 11(1):7:1–7:31
- [AMO13] Arthan R, Martin U, Oliva P (2013) A Hoare logic for linear systems. *Formal Asp Comput* 25(3):345–363
- [AO81] Apt KR, Olderog E-R (1981) Proof rules dealing with fairness. In: Logic of programs, volume 131 of Lecture notes in computer science. Springer, pp 1–8

- [AO83] Apt KR, Olderog E-R (1983) Proof rules and transformations dealing with fairness. *Sci Comput Program* 3:65–100
- [AO91] Apt KR, Olderog E-R (1991) Verification of sequential and concurrent programs. Springer, New York
- [AP86] Apt KR, Plotkin GD (1986) Countable nondeterminism and random assignment. *J ACM* 33(4):724–767
- [APS84] Apt KR, Pnueli A, Stavi J (1984) Fair termination revisited with delay. *Theor Comput Sci* 33:65–84
- [Apt81] Apt KR (1981) Ten years of Hoare's logic, a survey, part I. *ACM Trans Program Lang Syst* 3:431–483
- [Apt83] Apt KR (1983) Formal justification of a proof system for communicating sequential processes. *J ACM* 30:197–216
- [Apt84] Apt KR (1984) Ten years of Hoare's logic, a survey, part II: nondeterminism. *Theor Comput Sci* 28:83–109
- [Apt86] Apt KR (1986) Correctness proofs of distributed termination algorithms. *ACM Trans Program Lang Syst* 8:388–405
- [Bac80] Back R-JR (1980) Correctness preserving program refinements: proof theory and applications. Technical report 131, Mathematisch Centrum, Amsterdam
- [Bac86] Backhouse RC (1986) Program construction and verification. Prentice-Hall International, Englewood Cliffs
- [Bal06] Balser M (2006) Verifying concurrent systems with symbolic execution—temporal reasoning is symbolic execution with a little induction. PhD thesis, University of Augsburg. Shaker Verlag
- [BCD⁺05] Barnett M, Chang BE, DeLine R, Jacobs B, Leino KRM (2005) Boogie: a modular reusable verifier for object-oriented programs. In: de Boer FS, Bonsangue MM, Graf S, de Roeper WP (eds) Formal methods for components and objects, 4th international symposium, FMCO 2005, revised lectures, volume 4111 of Lecture notes in computer science. Springer, pp 364–387
- [Ben04] Benton N (2004) Simple relational correctness proofs for static analyses and program transformations. In: Jones ND, Leroy X (eds) Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on principles of programming languages (POPL 2004). ACM, pp 14–25
- [Ben12] Ben-Ari M (2012) Mathematical logic for computer science, 3rd edn. Springer
- [BG87] Blass A, Gurevich Y (1987) Existential fixed-point logic. In: Börger E (ed) Computation theory and logic. Springer, pp 20–36
- [BGA⁺14] Barthe G, Gaboardi M, Arias EJG, Hsu J, Kunz C, Strub P (2014) Proving differential privacy in Hoare logic. In: IEEE 27th computer security foundations symposium (CSF 2014). IEEE Computer Society, pp 411–424
- [BGB09] Barthe G, Grégoire B, Béguelin SZ (2009) Formal certification of code-based cryptographic proofs. In: Shao Z, Pierce BC (eds) Proceedings of the 36th ACM SIGPLAN-SIGACT symposium on principles of programming languages, POPL 2009. ACM, pp 90–101
- [BH16] Bubel R, Hähnle R (2016) KeY-Hoare. In: Deductive software verification—the KeY book—from theory to practice, volume 10001 of Lecture notes in computer science. Springer, pp 571–589
- [BKOB13] Barthe G, Köpf B, Olmedo F, Béguelin SZ (2013) Probabilistic relational reasoning for differential privacy. *ACM Trans Program Lang Syst* 35(3):9:1–9:49
- [BPS01] Bergstra JA, Ponse A, Smolka SA (eds) (2001) Handbook of process algebra. North-Holland/Elsevier
- [BR16] Baltag A, Renne B (2016) Dynamic epistemic logic. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/dynamic-epistemic>.
- [Bro07] Brookes S (2007) A semantics for concurrent separation logic. *Theor Comput Sci* 375(1–3):227–270
- [BRS⁺00] Balser M, Reif W, Schellhorn G, Stenzel K, Thums A (2000) Formal system development in KIV. In: Maibaum T (ed) Proceedings fundamental approaches to software engineering, volume 1783 of Lecture notes in computer science. Springer, pp 363–366
- [BT82a] Bergstra JA, Tucker JV (1982) Some natural structures which fail to possess a sound and decidable Hoare-like logic for their while-programs. *Theor Comput Sci* 17:303–315
- [BT82b] Bergstra JA, Tucker JV (1982) Expressiveness and the completeness of Hoare's logic. *J Comput Syst Sci* 25(3):267–284
- [BvW08] Back R-J, von Wright J (2008) Refinement calculus: a systematic introduction. Springer, New York
- [CDE⁺16] Certezeanu R, Drossopoulou S, Egelund-Müller B, Leino KRM, Sivarajan S, Wheelhouse MJ (2016) Quicksort revisited—verifying alternative versions of quicksort. In: Abraham E, Bonsangue MM, Johnsen EB (eds) Theory and practice of formal methods—essays dedicated to frank de boer on the occasion of his 60th birthday, volume 9660 of Lecture notes in computer science. Springer, pp 407–426
- [CGH83] Clarke EM, German SM, Halpern JY (1983) Effective axiomatizations of Hoare logics. *J ACM* 30(3):612–636
- [CGH⁺93] Clarke EM, Grumberg O, Hiraishi H, Jha S, Long DE, McMillan KL, Ness LA (1993) Verification of the futurebus+ cache coherence protocol. In: Agnew D, Claesen LJM, Camposano R (eds) Computer hardware description languages and their applications (CHDL '93), volume A-32 of IFIP transactions. North-Holland, pp 15–30
- [CGJ⁺03] Clarke E, Grumberg O, Jha S, Lu Y, Veith H (2003) Counterexample-guided abstraction refinement for symbolic model checking. *J ACM* 50(5):752–794
- [CH72] Clint M, Hoare CAR (1972) Program proving: jumps and functions. *Acta Inform* 1:214–224
- [Chl13] Chlipala A (2013) Certified programming with dependent types—a pragmatic introduction to the Coq Proof Assistant. MIT Press
- [CHVB18] Clarke EM, Henzinger TA, Veith H, Bloem R (eds) (2018) Handbook of model checking. Springer
- [Cla76] Clarke EM (1976) Completeness and incompleteness theorems for Hoare-like axiom systems. PhD thesis, Computer Science Department, Cornell University USA
- [Cla79] Clarke EM (1979) Programming language constructs for which it is impossible to obtain good Hoare axiom systems. *J ACM* 26(1):129–147
- [Cla85] Clarke EM (1985) The characterization problem for Hoare logics. In: Hoare CAR, Shepherdson JC (eds) Mathematical logic and programming languages. Prentice-Hall International, Englewood Cliffs, pp 89–106
- [Cli73] Clint M (1973) Program proving: coroutines. *Acta Inform* 2:50–63
- [CO81] Cartwright R, Oppen DC (1981) The logic of aliasing. *Acta Inform* 15:365–384
- [Coo78] Cook SA (1978) Soundness and completeness of an axiom system for program verification. *SIAM J Comput* 7(1):70–90
- [Coo81] Cook SA (1981) Corrigendum: soundness and completeness of an axiom system for program verification. *SIAM J Comput*, 10(3):612

- [Cou90] Cousot P (1990) Methods and logics for proving programs. In: Handbook of theoretical computer science, volume B: formal models and semantics (B). Elsevier, pp 841–994
- [CPR11] Cook B, Podelski A, Rybalchenko A (2011) Proving program termination. *Commun ACM* 54(5):88–98
- [dB75] de Bakker JW (1975) Inleiding Bewijsmethoden. In: Colloquium programmcorrectheid, MC Syllabus 21. Mathematisch Centrum, Amsterdam, pp 3–17
- [dB80] de Bakker JW (1980) Mathematical theory of program correctness. Prentice-Hall International, Englewood Cliffs
- [dB91] de Boer FS (1991) A compositional proof system for dynamic process creation. In: LICS. IEEE Computer Society, pp 399–405
- [dB99] de Boer FS (1999) A WP-calculus for OO. In: FoSSaCS, volume 1578 of Lecture notes in computer science. Springer, pp 135–149
- [dBP03] de Boer FS, Pierik C (2003) How to cook a complete Hoare logic for your pet OO language. In: FMCO, volume 3188 of Lecture notes in computer science. Springer, pp 111–133
- [DDH72] Dahl OJ, Dijkstra EW, Hoare CAR (eds) (1972). Structured programming. Academic Press Ltd.
- [DF88] Dijkstra EW, Feijen WHJ (1988) A method of programming. Addison-Wesley
- [dGdB⁺19] de Gouw S, de Boer FS, Bubel R, Hähnle R, Rot J, Steinhöfel D (2019) Verifying openjdk’s sort method for generic collections. *J Autom Reason* 62(1):93–126
- [dGR16] de Gouw S, Rot J (2016) Effectively eliminating auxiliaries. In: Ábrahám E, Bonsangue MM, Johnsen EB (eds) Theory and practice of formal methods—essays dedicated to frank de boer on the occasion of his 60th birthday, volume 9660 of Lecture notes in computer science. Springer, pp 226–241
- [dHdV02] den Hartog J, de Vink EP (2002) Verifying probabilistic programs using a Hoare like logic. *Int J Found Comput Sci* 13(3):315–340
- [Dij68] Dijkstra EW (1968) Cooperating sequential processes. In: Genuys F (ed) Programming languages: NATO advanced study institute. Academic Press, London, pp 43–112
- [Dij75] Dijkstra EW (1975) Guarded commands, nondeterminacy and formal derivation of programs. *Commun ACM* 18:453–457
- [Dij76a] Dijkstra EW (1976) A discipline of programming. Prentice-Hall, Englewood Cliffs
- [Dij76b] Dijkstra EW (1976) A great improvement. <http://www.cs.utexas.edu/users/EWD/ewd05xx/EWD573.PDF>, published as [Dij82]
- [Dij82] Dijkstra EW (1982) A great improvement. In: Selected writings on computing: a personal perspective. Springer, pp 217–219
- [DJ83] Damm W, Josko B (1983) A sound and relatively complete Hoare-logic for a language with higher type procedures. *Acta Inform* 20:59–101
- [dMB08] de Moura LM, Bjørner N (2008) Z3: an efficient SMT solver. In: Ramakrishnan CR, Rehof J (eds) Tools and algorithms for the construction and analysis of systems, TACAS 2008, volume 4963 of Lecture notes in computer science. Springer, pp 337–340
- [dRdBH⁺01] de Roever WP, de Boer FS, Hannemann U, Hooman J, Lakhnech Y, Poel M, Zwiers J (2001) Concurrency verification: introduction to compositional and noncompositional methods, volume 54 of Cambridge tracts in theoretical computer science. Cambridge University Press
- [DS90] Dijkstra EW, Scholten CS (1990) Predicate calculus and program semantics. Springer, New York
- [EC82] Emerson EA, Clarke EM (1982) Using branching time temporal logic to synthesize synchronization skeletons. *Sci Comput Program* 2(3):241–266
- [Eng17] Engelmann B (2017) Techniques for the verification of dynamically typed programs. PhD thesis, University of Oldenburg, Germany
- [EO16] Engelmann B, Olderog E-R (2016) A sound and complete Hoare logic for dynamically-typed, object-oriented programs. In: Ábrahám E, Bonsangue MM, Johnsen EB (eds) Theory and practice of formal methods—essays dedicated to frank de boer on the occasion of his 60th birthday, volume 9660 of Lecture notes in computer science. Springer, pp 173–193
- [FH71] Foley M, Hoare CAR (1971) Proof of a recursive program: quicksort. *Comput J* 14(4):391–395
- [Fil07] Filliâtre J-C (2007) Formal proof of a program: find. *Sci Comput Program* 64(3):332–340
- [Flo67] Floyd R (1967) Assigning meaning to programs. In: Schwartz JT (ed) Proceedings of symposium on applied mathematics 19, mathematical aspects of computer science. American Mathematical Society, New York, pp 19–32
- [FMV14] Furia CA, Meyer B, Velder S (2014) Loop invariants: analysis, classification, and examples. *ACM Comput Surv* 46(3):34:1–34:51
- [Fra86] Francez N (1986) Fairness. Springer, New York
- [Fra92] Francez N (1992) Program verification. Addison-Wesley, Reading
- [GCH89] German SM, Clarke EM, Halpern JY (1989) Reasoning about procedures as parameters in the language L4. *Inf Comput* 83(3):265–359
- [GFK84] Grumberg O, Francez N, Katz S (1984) Fair termination of communicating processes. In: Proceedings of the third annual ACM symposium on principles of distributed computing, Vancouver, B.C., Canada, August 27–29, 1984. ACM, pp 254–265
- [GFMDr81] Grumberg O, Francez N, Makowsky JA, de Roever WP (1981) A proof rule for fair termination of guarded commands. In: de Bakker J, van Vliet J (eds) Proceedings of the international symposium on algorithmic languages, pp 339–416
- [GFMDr85] Grumberg O, Francez N, Makowsky JA, de Roever WP (1985) A proof rule for fair termination of guarded commands. *Inf Control* 66(1/2):83–102
- [GL80] Gries D, Levin G (1980) Assignment and procedure call proof rules. *ACM Trans Program Lang Syst* 2(4):564–579
- [GMW79] Gordon MJC, Milner R, Wadsworth CP (1979) Edinburgh LCF, volume 78 of Lecture notes in computer science. Springer
- [Gor75] Gorelick GA (1975) A complete axiomatic system for proving assertions about recursive and nonrecursive programs. Technical report 75, Department of Computer Science, University of Toronto. <https://archive.org/details/ACompleteAxiomaticSystemForProvingAssertionsAboutRecursiveAnd>.
- [Gor88] Gordon MJC (1988) Programming language theory and its implementation—applicative and imperative paradigms. Prentice Hall International series in Computer Science. Prentice Hall
- [Gri78] Gries D (March 1978) The multiple assignment statement. *IEEE Trans Softw Eng* SE-4:89–93
- [Gri81] Gries D (1981) The science of programming. Springer, New York

- [Har79] Harel D (1979) First-order dynamic logic. Lecture notes in computer science 68. Springer, New York
- [HdR86] Hooman J, de Roever WP (1986) The quest goes on: a survey of proofsystems for partial correctness of CSP. In: Current trends in concurrency. Lecture notes in computer science 224. Springer, New York, pp 343–395
- [HJ00] Huisman M, Jacobs B (2000) Java program verification via a Hoare logic with abrupt termination. In: Fundamental approaches to software engineering, third international conference, FASE 2000, volume 1783 of Lecture notes in computer science. Springer, pp 284–303
- [HKT00] Harel D, Kozen D, Tiuryn J (2000) Dynamic logic. MIT Press, Cambridge
- [HL74] Hoare CAR, Lauer PE (1974) Consistent and complementary formal theories of the semantics of programming languages. Acta Inform 3:135–153
- [HN18] Haslbeck MPL, Nipkow T (2018) Hoare logics for time bounds. Archive of formal proofs. http://isa-afp.org/entries/Hoare_Time.html. Formal proof development.
- [Hoa61] Hoare CAR (1961) Algorithm 64: quicksort. Commun ACM 4(7):321
- [Hoa69] Hoare CAR (1969) An axiomatic basis for computer programming. Commun ACM 12:576–580, 583
- [Hoa71a] Hoare CAR (1971) Procedures and parameters: an axiomatic approach. In: Engeler E (ed) Proceedings of symposium on the semantics of algorithmic languages. Lecture notes in mathematics 188. Springer, New York, pp 102–116
- [Hoa71b] Hoare CAR (1971) Proof of a program: FIND. Commun ACM 14(1):39–45
- [Hoa72a] Hoare CAR (1972) An axiomatic definition of the programming language PASCAL. In: International symposium on theoretical programming, volume 5 of Lecture notes in computer science. Springer, pp 1–16
- [Hoa72b] Hoare CAR (1972) Proof of a structured program: 'the sieve of Eratosthenes'. Comput J 15(4):321–325
- [Hoa72c] Hoare CAR (1972) Towards a theory of parallel programming. In: Hoare CAR, Perrot RH (eds) Operating systems techniques. Academic Press, London, pp 61–71
- [Hoa75] Hoare CAR (1975) Parallel programming: an axiomatic approach. Comput Lang 1(2):151–160
- [Hoa78] Hoare CAR (1978) Communicating sequential processes. Commun ACM 21:666–677
- [HOP10] Hoenicke J, Olderog E-R, Podelski A (2010) Fairness for dynamic control. In: Esparza J, Majumdar R (eds) Tools and algorithms for the construction and analysis of systems, 16th international conference, (TACAS 2010), volume 6015 of Lecture notes in computer science. Springer, pp 251–265
- [HP15] Hoenicke J, Podelski A (2015) Fairness for infinitary control. In: Meyer R, Platzer A, Wehrheim H (eds) Correct system design—symposium in Honor of Ernst-Rüdiger olderog on the occasion of his 60th birthday, volume 9360 of Lecture notes in computer science. Springer, pp 33–43
- [HW73] Hoare CAR, Wirth N (1973) An axiomatic definition of the programming language PASCAL. Acta Inform 2:335–355
- [INM84] INMOS Limited (1984) Occam programming manual. Prentice-Hall International, Englewood Cliffs
- [Jon81] Jones CB (1981) Developing methods for computer programs including a notion of interference. PhD thesis, University of Oxford, UK
- [Jon83] Jones CB (1983) Tentative steps toward a development method for interfering programs. ACM Trans Program Lang Syst 5(4):596–619
- [Jon03] Jones CB (2003) The early search for tractable ways of reasoning about programs. IEEE Ann Hist Comput 25(2):26–49
- [JP04] Jacobs B, Poll E (2004) Java program verification at nijmegen: developments and perspective. In: Futatsugi K, Mizoguchi F, Yonezaki N (eds) Software security—theories and systems, second Mext-NSF-JSPS international symposium, ISSS 2003, Tokyo, Japan, November 4–6, 2003, revised papers, volume 3233 of Lecture notes in computer science. Springer, pp 134–153
- [JR10] Jones CB, Roscoe AW (2010) Insight, inspiration and collaboration. In: Roscoe AW, Jones CB, Wood KR (eds) Reflections on the work of C.A.R. Hoare. Springer, pp 1–32
- [JW75] Jensen K, Wirth N (1975) PASCAL user manual and report. Springer
- [Kal90] Kaldewaij A (1990) Programming: the derivation of algorithms. Prentice-Hall International, Englewood Cliffs
- [KGJ+15] Katoen J, Gretz F, Jansen N, Kaminski BL, Olmedo F (2015) Understanding probabilistic programs. In: Meyer R, Platzer A, Wehrheim H (eds) Correct system design—symposium in Honor of Ernst-Rüdiger olderog on the occasion of his 60th birthday, volume 9360 of Lecture notes in computer science. Springer, pp 15–32
- [Kin69] King J (1969) Developing methods for computer programs including a notion of interference. PhD thesis, Department of Computer Science, Carnegie-Mellon University, Pittsburgh USA
- [Kle98] Kleymann T (1998) Hoare logic and VDM : machine-checked soundness and completeness proofs. PhD thesis, University of Edinburgh UK
- [Kle99] Kleymann T (1999) Hoare logic and auxiliary variables. Formal Asp Comput 11(5):541–566
- [Koz00] Kozen D (2000) On Hoare logic and Kleene algebra with tests. ACM Trans Comput Log 1(1):60–76
- [KT01] Kozen D, Tiuryn J (2001) On the completeness of propositional Hoare logic. Inf Sci 139(3-4):187–195
- [Lam77] Lamport L (1977) Proving the correctness of multiprocess programs. IEEE Trans Softw Eng SE-3:2:125–143
- [Lam94] Lamport L (1994) The temporal logic of actions. ACM Trans Program Lang Syst 16(3):872–923
- [Lam02] Lamport L (2002) Specifying systems, the TLA+ language and tools for hardware and software engineers. Addison-Wesley
- [Lan79] Langmaack H (1979) A proof of a theorem of Lipton on Hoare logic and applications. Technical report, Ber. 8003, Inst Inf Prakt Math, University of Kiel, Germany
- [Lan82] Langmaack H (1982) On the termination problem for finitely interpreted ALGOL-like programs. Acta Inform 18:79–108
- [Lau71] Lauer PE (1971) Consistent formal theories of the semantics of programming languages. Technical report 25. 121, IBM Laboratory Vienna
- [LCC+05] Leavens GT, Cheon Y, Clifton C, Ruby C, Cok DR (2005) How the design of JML accommodates both runtime assertion checking and formal verification. Sci Comput Program 55(1-3):185–208
- [Lei10] Leino KRM (2010) Dafny: an automatic program verifier for functional correctness. In: Clarke EM, Voronkov A (eds) Logic for programming, artificial intelligence, and reasoning, LPAR-16, volume 6355 of Lecture notes in computer science. Springer, pp 348–370
- [LG81] Levin G, Gries D (1981) A proof technique for communicating sequential processes. Acta Inform 15:281–302

- [LGH⁺78] London RL, Gutttag JV, Horning JJ, Lampson BW, Mitchell JG, Popek GJ (1978) Proof rules for the programming language Euclid. *Acta Inform* 10:1–26
- [Lip75] Lipton RJ (1975) Reduction: a method of proving properties of parallel programs. *Commun ACM* 18:717–721
- [Lip77] Lipton RJ (1977) A necessary and sufficient condition for the existence of Hoare logics. In: 18th annual symposium on foundations of computer science. IEEE Computer Society, pp 1–6
- [LO80] Langmaack H, Olderog E-R (1980) Present-day Hoare-like systems for programming languages with procedures: power, limits and most likely expressions. In: de Bakker JW, van Leeuwen J (eds) *Automata, languages and programming*, 7th colloquium, volume 85 of *Lecture notes in computer science*. Springer, pp 363–373
- [LPS81] Lehmann DJ, Pnueli A, Stavi J (1981) Impartiality, justice, and fairness: the ethics of concurrent termination. In: Kariv O, Even S (eds) *Proceedings of international colloquium on automata languages and programming (ICALP '81)*. Lecture notes in computer science 115. Springer, New York, pp 264–277
- [LS87] Loeckx J, Sieber K (1987) *The foundation of program verification*, 2nd edn. Teubner-Wiley, Stuttgart
- [Man74] Manna Z (1974) *Mathematical theory of computation*. Mc Graw-Hill
- [Mey97] Meyer B (1997) *Object-oriented software construction*, 2nd edn. Prentice Hall
- [Mil80] Milner R (1980) *A calculus of communicating systems*. Lecture notes in computer science 92, Springer, New York
- [Mil89] Milner R (1989) *Communication and concurrency*. Prentice-Hall International, Englewood Cliffs
- [MJ84] Morris FL, Jones CB (1984) An early program proof by Alan Turing. *Ann Hist Comput* 6:139–143
- [MM05] McIver A, Morgan C (2005) *Abstraction, refinement and proof for probabilistic systems*. Monographs in computer science. Springer
- [Moi83] Moitra A (1983) On Apt, Francez, and de Roever's "A proof system for communicating sequential processes". *ACM Trans Program Lang Syst* 5(3):500–501
- [Mor94] Morgan C (1994) *Programming from specifications*, 2nd edn. Prentice-Hall International, London
- [MP74] Manna Z, Pnueli A (1974) Axiomatic approach to total correctness of programs. *Acta Inform* 3:253–263
- [MP91] Manna Z, Pnueli A (1991) *The temporal logic of reactive and concurrent systems—specification*. Springer, New York
- [MP95] Manna Z, Pnueli A (1995) *Temporal verification of reactive systems—safety*. Springer, New York
- [MS87] Mirkowska C, Salwicki A (1987) *Algorithmic logic*. Kluwer Academic Publishers, Norwell
- [Nau66] Naur P (1966) Proof of algorithms by general snapshots. *BIT Numer Math* 6(4):310–316
- [NBB⁺63] Naur P, Backus J, Bauer F, Green J, Katz C, McCarthy J, Perlis A, Rutishauser H, Samelson K, Vauquois B, Wegstein J, van Wijngaarden A, Woodger M (1963) Report on the algorithmic language ALGOL 60. *Numer Math* 4:420–453
- [NCMM09] Nordio M, Calcagno C, Müller P, Meyer B (2009) A sound and complete program logic for eiffel. In: Oriol M, Meyer B (eds) *Objects, components, models and patterns*, 47th international conference, TOOLS EUROPE 2009, Zurich, Switzerland, June 29–July 3, 2009. *Proceedings*, volume 33 of *Lecture notes in business information processing*. Springer, pp 195–214
- [Nie87] Nielson HR (1987) A Hoare-like proof system for analysing the computation time of programs. *Sci Comput Program* 9(2):107–136
- [Nip02a] Nipkow T (2002) Hoare logics in Isabelle/HOL. In: Schwichtenberg H, Steinbrüggen R (eds) *Proof and system-reliability*, volume 62 of *NATO science series*. Springer, pp 341–367
- [Nip02b] Nipkow T (2002) Hoare logics for recursive procedures and unbounded nondeterminism. In: *Proceedings of the computer science logic*, 16th international workshop, CSL 2002, volume 2471 of *Lecture notes in computer science*. Springer, pp 103–119
- [NN99] Nipkow T, Nieto LP (1999) Owicki/Gries in Isabelle/HOL. In: Finance JP (ed) *Fundamental approaches in software engineering (FASE)*, volume 1577 of *Lecture notes in computer science*. Springer, pp 188–203
- [NPW02] Nipkow T, Paulson LC, Wenzel M (2002) Isabelle/HOL—A proof assistant for higher-order logic, volume 2283 of *Lecture notes in computer science*. Springer
- [OA88] Olderog E-R, Apt KR (1988) Fairness in parallel programs, the transformational approach. *ACM Trans Program Lang Syst* 10:420–455
- [OG76a] Owicki S, Gries D (1976) An axiomatic proof technique for parallel programs. *Acta Inform* 6:319–340
- [OG76b] Owicki S, Gries D (1976) Verifying properties of parallel programs: an axiomatic approach. *Commun ACM* 19:279–285
- [O'H07] O'Hearn PW (2007) Resources, concurrency, and local reasoning. *Theor Comput Sci* 375(1–3):271–307
- [O'H19] O'Hearn PW (2019) Separation logic. *Commun ACM* 62(2):86–95
- [Old81] Olderog E-R (1981) Sound and complete Hoare-like calculi based on copy rules. *Acta Inform* 16:161–197
- [Old83a] Olderog E-R (1983) A characterization of Hoare's logic for programs with Pascal-like procedures. In: *Proceedings of the 15th ACM symposium on theory of computing (STOC)*. ACM, pp 320–329
- [Old83b] Olderog E-R (1983) On the notion of expressiveness and the rule of adaptation. *Theor Comput Sci* 30:337–347
- [Old84] Olderog E-R (1984) Correctness of programs with Pascal-like procedures without global variables. *Theor Comput Sci* 30:49–90
- [OP10] Olderog E-R, Podelski A (2010) Explicit fair scheduling for dynamic control. In: Dams D, Hannemann U, Steffen M (eds) *Concurrency, compositionality, and correctness, essays in Honor of Willem-Paul de Roever*, volume 5930 of *Lecture notes in computer science*. Springer, pp 96–117
- [ORS92] Owre S, Rushby JM, Shankar N (1992) PVS: A prototype verification system. In: Kapur D (ed) *Automated deduction (CADE-11)*, 11th international conference on automated deduction, 1992, proceedings, volume 607 of *Lecture notes in computer science*. Springer, pp 748–752
- [ORY01] O'Hearn PW, Reynolds JC, Yang H (2001) Local reasoning about programs that alter data structures. In: Fribourg L (ed) *Computer science logic*, 15th international workshop (CSL 2001), volume 2142 of *Lecture notes in computer science*. Springer, pp 1–19
- [OW12] Olderog E-R, Wilhelm R (2012) Turing und die Verifikation. *Informatik Spektrum* 35(4):271–279
- [Owi75] Owicki S (1975) *Axiomatic proof techniques for parallel programs*. Outstanding dissertations in the computer sciences. Garland Publishing, New York
- [Owi76] Owicki S (1976) A consistent and complete deductive system for the verification of parallel programs. In: *STOC*. ACM, pp 73–86

- [Owi78] Owicki S (1978) Verifying concurrent programs with shared data classes. In: Neuhold EJ (ed) Proceedings of the IFIP working conference on formal description of programming concepts. North-Holland, Amsterdam, pp 279–298
- [OYR04] O'Hearn PW, Yang H, Reynolds JC (2004) Separation and information hiding. In: Jones ND, Leroy X (eds) Proceedings of 31st symposium on principles of programming languages (POPL 2004). ACM, pp 268–280
- [PdB03] Pierik C, de Boer FS (2003) A syntax-directed Hoare logic for object-oriented programming concepts. In: FMOODS, volume 2884 of Lecture notes in computer science. Springer, pp 64–78
- [Pla08] Platzer A (2008) Differential dynamic logic for hybrid systems. *J Autom Reasoning* 41(2):143–189
- [Pnu77] Pnueli A (1977) The temporal logic of programs. In: Proceedings of the 18th IEEE symposium on foundations of computer science, pp 46–57
- [PR18] Piskac R, Rümmer P (eds) (2018) Verified software. Theories, tools, and experiments—10th international conference, VSTTE 2018, Oxford, UK, July 18–19, 2018, Revised selected papers, volume 11294 of Lecture notes in computer science. Springer
- [Pra76] Pratt VR (1976) Semantical considerations on Floyd-Hoare logic. In: 17th annual symposium on foundations of computer science (FoCS 1976). IEEE Computer Society, pp 109–121
- [QS81] Queille JP, Sifakis J (1981) Specification and verification of concurrent systems in CESAR. In: Proceedings of the 5th international symposium on programming, Paris
- [Rey02] Reynolds JC (2002) Separation logic: a logic for shared mutable data structures. In: 17th IEEE symposium on logic in computer science (LICS 2002). IEEE Computer Society, pp 55–74
- [Sal70] Salwicki A (1970) Formalized algorithmic languages. *Bull Acad Pol Sci* 18:227–232
- [Sch77] Schwarz J (1977) Generic commands—a tool for partial correctness formalisms. *Comput J* 20:151–155
- [Sok77] Sokołowski S (1977) Total correctness for procedures. In: 6th symposium on mathematical foundations of computer science, volume 53 of Lecture notes in computer science. Springer, pp 475–483
- [Sok87] S. Sokołowski (1987) Soundness of Hoare's logic: an automated proof using LCF. *ACM Trans Program Lang Syst* 9(1):100–120
- [Spi92] Spivey JM (1992) The Z notation: a reference manual, 2nd edn. Prentice Hall
- [SS14] Sznuk T, Schubert A (2014) Tool support for teaching Hoare logic. In: Giannakopoulou D, Salaün G (eds) Software engineering and formal methods—12th international conference (SEFM 2014), volume 8702 of Lecture notes in computer science. Springer, pp 332–346
- [Tar36] Tarski A (1936) Der Wahrheitsbegriff in den formalisierten Sprachen. *Studia Philosophica* 1(3):261–405
- [Tur49] Turing AM (1949) On checking a large routine. Report of a conference on high speed automatic calculating machines, pp 67–69, 1949. University Mathematics Laboratory, Cambridge. (See also: F. L. Morris and Jones CB, *An early program proof by Alan Turing*, *Annals of the History of Computing* 6 pp 139–143, 1984)
- [TZ88] Tucker JV, Zucker JI (1988) Program correctness over abstract data types, with error-state semantics. North-Holland and CWI Monographs, Amsterdam
- [Unr19] Unruh D (2019) Quantum relational Hoare logic. *Proc ACM Program Lang* 3(POPL):33:1–33:31
- [vO99] von Oheimb D (1999) Hoare logic for mutual recursion and local variables. In: Foundations of software technology and theoretical computer science, 19th conference, Chennai, India, December 13–15, 1999, proceedings, volume 1738 of Lecture notes in computer science. Springer, pp 168–180
- [vON02] von Oheimb D, Nipkow T (2002) Hoare logic for nanojava: auxiliary variables, side effects, and virtual methods revisited. In: Eriksson L, Lindsay PA (eds) FME 2002: formal methods—getting IT right, international symposium of formal methods Europe, Copenhagen, Denmark, July 22–24, 2002, proceedings, volume 2391 of Lecture notes in computer science. Springer, pp 89–105
- [vWMP⁺75] van Wijngaarden A, Mailloux BJ, Peck JEL, Koster CHA, Sintzoff M, Lindsey CH, Meertens LGLT, Fisker RG (1975) Revised report on the algorithmic language ALGOL 68. *Acta Inform* 5:1–236
- [Wan78] Wand M (1978) A new incompleteness result for Hoare's system. *J ACM* 25(1):168–175
- [WH66] Wirth N, Hoare CAR (1966) A contribution to the development of ALGOL. *Commun ACM* 9(6):413–432
- [Win93] Winskel G (1993) The formal semantics of programming languages—an introduction. Foundation of computing series. MIT Press
- [Yin11] Ying M (2011) Floyd–Hoare logic for quantum programs. *ACM Trans Program Lang Syst* 33(6):19:1–19:49
- [Yin19] Ying M (2019) Toward automatic verification of quantum programs. *Formal Asp Comput* 31(1):3–25
- [Zöb88] Zöbel D (1988) Normalform-Transformationen für CSP-Programme. *Informatik: Forschung und Entwicklung* 3:64–76

Received 8 April 2019

Accepted in revised form 8 November 2019 by Cliff Jones and José N. Oliveira

Published online 27 November 2019