

Verification Conditions for Source-level Imperative Programs

Maria João Frade and Jorge Sousa Pinto

*Departamento de Informática / CCTC
Universidade do Minho, Braga, Portugal*

Abstract

This paper is a systematic study of verification conditions and their use in the context of program verification. We take Hoare logic as a starting point and study in detail how a verification conditions generator can be obtained from it. The notion of *program annotation* is essential in this process. Weakest preconditions and the use of *updates* are also studied as alternative approaches to verification conditions. Our study is carried on in the context of a While language. Important extensions to this language are considered toward the end of the paper. We also briefly survey modern program verification tools and their approaches to the generation of verification conditions.

Key words: Hoare logic, Verification Conditions, Program Verification, Program Annotations, Weakest Preconditions, Updates

1 Introduction

The idea of specifying the behaviour of programs through the use of preconditions and postconditions, and in general of assertions that are true or false relative to the current state of execution of a program, has been around since the 1960's and given rise to the development of the *axiomatic* style of program semantics, associated with the use of some *program logic*. The same idea has been used since the 1970's in the implementation of practical tools for checking the correct behaviour of programs (vis-à-vis their specifications), marking the birth of program verification as a research area.

Email address: {mjf,jsp}@di.uminho.pt (Maria João Frade and Jorge Sousa Pinto).

In the development of early tools [49, 40, 58, 15] it became clear that the most convenient way to organise a program verification system is to use the axiomatic semantics to generate first-order proof obligations (baptised *Verification Conditions*, VCs for short) that can be handled by a standard proof tool. The idea is that if all proof obligations generated for a program can be discharged (i.e. they can be proved valid), then the program is guaranteed to be correct. The semantics-based component that generates the proof obligations is called a *Verification Conditions Generator* (VCGen for short).

Program verification has recently received renewed attention from the software engineering community. One very general reason for this is the continuing and increasing pressure on industry to deliver software that can be certified as safe and correct. A more specific reason is that program verification methods suit very naturally the so-called *design-by-contract* methodology for software development, with the advent of program annotation languages like JML. Program verification fits in as the static, formal component of a methodology that encompasses also other validation methods like dynamic checking and testing.

As will become clear throughout the paper, program verification based on program logics cannot in general be fully automated. Human intervention is often required in one (or both) of two forms: by providing *annotations* in the code, such as loop invariants, that facilitate automatic reasoning, and by manually proving the hardest proof obligations generated in the verification process – the need for manual proof (usually with the help of a proof assistant) is in fact a direct consequence of the fact that first-order logic is not decidable.

We remark however that there exists a wealth of work on the automatic generation of invariants (see Section 11 for classic references), and some modern verification platforms incorporate invariant generation functionality. Moreover, advances in automated theorem proving have also been very significant in recent years. Another reason for the recent resurgence of program verification as a hot topic has to do precisely with advances in both of these areas, in particular the emergence of SMT solver technology. We shall have more to say about this in Section 11.

In contrast with this flavour of program verification (usually known as *deductive* verification), we find a class of techniques in which automation is primary. Techniques organised around the designation of *software model checking* are based on the extraction of a labelled transition system from the code, followed by the verification of temporal properties by model checking the transition system (i.e. by exploring the model exhaustively). Several surveys are available on software model checking [32, 30, 53]. The first of these papers classifies deductive methods as *path insensitive*, and software-model checking methods as *path sensitive*. This means in particular that when a problem is

found, a software model checker returns one particular initial state that leads to that problem, whereas deductive methods consider sets of initial states, characterised by first-order properties. In the case of concurrent systems, all paths corresponding to different interleavings of the processes involved are considered, which makes path sensitive techniques particularly suitable to analyse such systems.

Full automation is of course a standard argument in favour of model checking techniques. Let us note however that automation comes with a price, which in the case of model checking is the well-known state explosion problem. In practice, software model checkers typically rely on additional techniques such as *predicate abstraction*, which allows for the construction of over-approximated models, with significantly smaller state spaces. Precision is lost, but soundness is not, which means that although false errors may be found, absence of errors can be interpreted safely.

It is also essential to mention the theoretical framework of *abstract interpretation* [22], which stands at the heart of many successful static analysers. Although not usually considered as a form of program verification, abstract interpretation is used as an auxiliary tool for both deductive verification (notably in the automatic inference of invariants, see Section 11) and software model checking tools (predicate abstraction, mentioned above, is a form of abstract interpretation).

In comparison with path-sensitive approaches [32], deductive methods benefit from very expressive property languages (typically based on first-order logic) and are particularly amenable to compositional reasoning – the design-by-contract approach explores precisely this aspect. A more technological advantage, which has been explored by a few recent tools, has to do with the generation of first-order proof obligations. Since there are many available proof tools, which are continuously being developed and improved, verification platforms can take advantage of those improvements, and even use a combination of proof tools, both automatic and interactive. The analysis of the source code and generation of proof obligations is completely decoupled from the first-order theorem proving.

Let us add to this that the flexibility of deductive methods has in recent times allowed these methods to be successfully used in targeting new certification requirements created, for instance, by mobile code and the associated new architectures for the execution of software. For instance, the following two techniques both rely on verification conditions, and provide yet further motivation for the relevance of program logic-based methods for reasoning about programs.

- *Proof-carrying code* [71], based on the generation of verification conditions

from annotated low-level (compiled) code. The idea is that a compiler can automatically produce a proof – a certificate – that the compiled code satisfies some requirements (say, it performs only safe memory accesses), and an execution platform may then (cheaply, and without relying on the behaviour of complex pieces of software like compilers or theorem provers) generate VCs and check that the certificate provides evidence for these VCs. The very existence of this certificate, a *proof object*, is a unique characteristic of deductive methods.

- Certain information-flow properties (such as non-interference) that have traditionally been treated using a number of language-based security techniques [78] may also be addressed using deductive methods like *self-composition* [8], which can be implemented resorting to a standard deductive verification platform based on Hoare logic or weakest preconditions. *Dynamic logic* has also been used to the same effect [23] but requires a dedicated tool.

This paper is a study of verification conditions for imperative, sequential, high-level programming languages. Our tutorial presentation treats annotated programs formally, and provides a uniform development of VCGen algorithms from program logics. Although we do not attempt to give a thorough survey of program logics or program verification systems, we do examine (and provide references for) those systems that have become more popular in recent years.

The paper is structured as follows. Section 2 sets the basis by introducing a simple programming language and the notion of Hoare triple. In Section 3 the inference system H of Hoare logic is introduced, and Section 4 discusses its use in program verification, based on the generation of Verification Conditions. Section 5 presents an alternative, goal-directed formulation of Hoare logic (system H_g), that is more amenable to mechanising the construction of proof trees, since it contains no ambiguity in the choice of inference rule. In Section 6 we show how the introduction of program annotations (resulting in systems H_{ga} and H_{gi}) completes this progression toward the mechanisation of Hoare logic.

The previous approach forces the insertion of annotations between any two composed commands of a program. The next step is to eliminate the need for this tedious process. This can be done in two ways: by backward propagation of assertions, and by forward propagation. The former approach leads us in Section 7 to VCGens based on Dijkstra’s weakest preconditions. The latter approach, covered in Section 9, takes us in the direction of systems based on the use of *updates* (system H_u). Section 10 discusses how the simple language used here can be extended with a number of features that can be found in realistic programming languages. Finally, Section 11 surveys work that has helped to bring deductive verification into practice.

We also briefly review verification condition generation tools that bring into practice these ideas. Section 8 presents a guarded command language in the style introduced by Dijkstra, and discusses how it has been used as an intermediate language in the VCGens of the ESC family of tools. The KeY tool, based on a variant of dynamic logic, is discussed in Section 9; the Boogie and Why tools, both generic VCGens, are discussed in Section 11. Our goal here is not to exhaustively cover all tools for program verification; we merely select a few that we find suitable to illustrate the concepts discussed in the paper.

2 Programs and Specifications: Hoare Triples

In this paper we explore methods for specifying programs in a simple imperative language and for proving formally the correctness of programs with respect to such specifications. We consider only *partial correctness* specifications: programs are required to behave properly if they terminate, but are not required to terminate.

We consider a typical While language with data types for integer numbers and Booleans. Commands include a do-nothing command, assignment, composition, while loop and (two-branched) conditional execution. The language has two base types

$$\tau ::= \mathbf{bool} \mid \mathbf{int}$$

The syntax of Boolean and integer expressions could be more or less evolved; Boolean expressions should contain the Boolean constants **true** and **false**, and it must be possible to form expressions by comparing the values of two integers. The language also includes Boolean operators for conjunction, disjunction, and negation (we use C/Java-like syntax for operators). Integer expressions are formed from constants and a set of variables \mathcal{V} , together with a number of operators on integers. We let x, y, \dots range over \mathcal{V} . Note that these choices have no impact on the material presented in the paper; the reader will find in the literature presentations that use simpler languages as well as richer ones.

The key semantic notion is that of the program *state* (given by the values of the variables involved in the computation). The value of an expression depends on the current state, and the computing device may in general change the state when it executes a command. In addition to expressions and commands, we need syntax for formulas that express assertions about properties of particular states, as well as a class of formulas for specifying the behaviour of programs. We have the following *phrase types*

$$\theta ::= \mathbf{Exp}[\tau] \mid \mathbf{Comm} \mid \mathbf{Assert} \mid \mathbf{Spec}$$

Exp[int]	$\ni e ::= \dots -1 0 1 \dots x -e e + e e - e e * e e \text{div} e e \text{mod} e$
Exp[bool]	$\ni b ::= \text{true} \text{false} e == e e < e e <= e e > e e >= e e != e$ $ b \&\& b b b !b$
Comm	$\ni C ::= \text{skip} C ; C x := e \text{if } b \text{ then } C \text{ else } C \text{while } b \text{ do } C$
Assert	$\ni A ::= \text{true} \text{false} e == e e < e e <= e e > e e >= e e != e$ $!A A \&\& A A A A \rightarrow A \text{Forall } x. A \text{Exists } x. A$
Spec	$\ni S ::= \{A\} C \{A\}$

Fig. 1. Abstract Syntax

corresponding respectively to *expressions* (for each data type), *commands*, *assertions* and *specifications*. Their abstract syntax is defined in Figure 1. Note that the syntax rules may also be read implicitly as typing rules. For instance, if e_1, e_2 have type **Exp[int]**, then so has $e_1 + e_2$, and so on.

Expressions and commands are fairly obvious. It remains to discuss the language of assertions, i.e. properties that hold at a given point in the program execution, and program specifications.

It is common for assertions to be defined as a super-set of Boolean expressions, since they may have to refer to the values of expressions in the current state of the program. If the syntax for assertions is compatible with that of Boolean expressions, it will be easier for ordinary programmers to write specifications (and also *program annotations*, see Section 6).

We thus construct our language of assertions starting from the language of Boolean expressions of the programming language, and extend that with an implication connective. Moreover, to allow for first-order reasoning about programs, universal and existential quantifiers are introduced. So basically our language of assertions is a first-order language, with integer expressions as terms and binary comparison operators as predicates, both directly inherited from the programming language. Integer expressions are interpreted in \mathbb{Z} in the standard way, as will be detailed below.

In examples to be given later in the paper, we will consider the possibility of extending the assertion language with functions and predicates that are not part of the programming language. This possibility is offered by most program verification systems, and allows one to express properties that would otherwise not be possible. In Section 6 an example will illustrate this point, by introducing a logical function corresponding to Fibonacci numbers.

We remark that the particular choice for the language of expressions in the programming language is itself merely illustrative; an alternative, more generic

presentation would be possible, in which the expression language, as well as assertion language and its interpretation structure, are not fixed (this approach is followed in [66]). The main ideas developed in this paper do not depend on any particular choice of expression language, but fixing this language allows us to give concrete examples.

Assertions that hold before and after execution of a program – *preconditions* and *postconditions* respectively – will allow one to write specifications of programs or *Hoare triples* – the last syntactic category in Figure 1. The intuitive meaning of a specification $\{P\} C \{Q\}$ is that if the program C is executed in an initial state in which the assertion (precondition) P is *true*, then either execution of C does not terminate or, if it does, assertion Q (a postcondition) will be *true* in the final state. Because termination is not guaranteed, this notion is called a *partial correctness* specification.

The notions of specification and Hoare triple coincide. We remark however that sometimes it is useful to consider a notion of specification (P, Q) consisting only of a precondition P and a postcondition Q . In this view, if the Hoare triple $\{P\} C \{Q\}$ is valid for some program C , then C is said to be correct with respect to the specification (P, Q) . We will use both notions in the paper but this should not generate any confusion.

The introduction of binding quantifiers in assertions imposes the usual notions of free and bound variables. Variables that occur in the program C (and possibly also free in the precondition P or in the postcondition Q) are called *program variables*. Variables that occur free in P or Q but not in the program will be called *auxiliary* or *ghost variables* (their use will be explained in Section 7), and those that are bound by some quantifier in P or Q will be called *logical variables*.

Semantics. The meaning of a grammatically correct program can be formalised in different ways (see [84, 76]):

- *Operational semantics* is focused on the computation the program induces on a machine (*small-step*, or *structural*, if the emphasis is on the individual steps of the execution; *big-step*, or *natural semantics*, if the emphasis is on the relationship between the initial and the final state of the execution).
- *Denotational semantics* is focused on representing the effect of executing a program by a mathematical object.

Frequently, the logical system for proving partial correctness properties of programs is viewed as an *axiomatic semantics*, focused on specific properties (expressed by assertions) of the effect of executing a program.

In the following, a natural semantics is used to describe the meaning of commands, and we define semantic functions to interpret expressions, assertions and specifications in a denotational style.

The semantics is given in terms of states. The base types are interpreted as expected

$$\begin{aligned} \llbracket \mathbf{bool} \rrbracket &= \{true, false\} \\ \llbracket \mathbf{int} \rrbracket &= \mathbb{Z} \end{aligned}$$

Boolean and integer expressions are interpreted as Boolean or integer values, but these values depend on the values of variables that may occur in the expressions. In other words, they depend on a *state*, which is a function that maps each variable into its integer value.¹ We write $\Sigma = \mathcal{V} \rightarrow \llbracket \mathbf{int} \rrbracket$ for the set of states, and for $s \in \Sigma$, $y \in \mathcal{V}$ and $v \in \llbracket \mathbf{int} \rrbracket$, $s[y \mapsto v]$ denotes the following state

$$s[y \mapsto v](x) = \begin{cases} v & \text{if } x = y \\ s(x) & \text{if } x \neq y \end{cases}$$

For each phrase type the corresponding *domain of interpretation* (the set of possible meanings) is given as follows

$$\begin{aligned} \llbracket \mathbf{Exp}[\tau] \rrbracket &= \Sigma \rightarrow \llbracket \tau \rrbracket \\ \llbracket \mathbf{Comm} \rrbracket &\subseteq \Sigma \times \Sigma \\ \llbracket \mathbf{Assert} \rrbracket &= \Sigma \rightarrow \{true, false\} \\ \llbracket \mathbf{Spec} \rrbracket &= \{true, false\} \end{aligned}$$

These domains reflect our assumption that an expression has a value at every state (evaluation of expressions always terminates without an error stop, see below) and that expression evaluation never changes the state (the language is free of *side effects*).

The behaviour of a command is to transform the state of a computation. Commands are interpreted operationally via the *evaluation* relation $(\cdot, \cdot) \Downarrow \cdot \subseteq \mathbf{Comm} \times \Sigma \times \Sigma$. Intuitively $(C, s) \Downarrow s'$ means that the execution of C from s will terminate and the resulting state will be s' . A command whose execution does not terminate for an initial state s is absent from this interpretation since there is no pair (s, s') corresponding to it.

Denotationally, the meaning of a command would be a state-transformation function $\Sigma \rightarrow \Sigma$, and dealing with non-termination would require the in-

¹ Another possibility would be to consider states as partial functions.

roduction of more sophisticated mathematical domains. This denotational interpretation is not required (nor the most appropriate) for our present goal.

Assertions are interpreted as truth values depending on a valuation function given by the state, and specifications are interpreted as truth values independently of states.

Figure 2 shows the semantic equations that define the interpretation functions for expressions, assertions and specifications, and the natural semantics for commands as a set of evaluation rules. Note that we make an overloaded use of $\llbracket \cdot \rrbracket$ (we could subscript the semantic brackets with the phrase type of the object phrase that is being interpreted, but this is usually obvious from the phrase itself or from the context).

Concerning the interpretation of expressions, note that we are assuming that $n \div 0$ and $n \bmod 0$ produce some erroneous (and fixed) integer result. To treat the detection of arithmetic errors it would be necessary to extend our current semantics, for instance by enlarging the domains of interpretation of integer and Boolean expressions to include one or more special results denoting errors.

Note that the semantic interpretation of a command C can be seen as a partial function, since the binary relation on states induced by C satisfies the following property.

Lemma 1 (Determinacy) *If $(C, s) \Downarrow s'$ and $(C, s) \Downarrow s''$, then $s' = s''$.*

Proof. By induction on the structure of C . □

The semantic interpretation of assertions is the usual for first-order logic. Observe the correspondence between the logical connectives of the assertion language and the connectives of classic predicate calculus used in the metalogic. Let A be an assertion and $s \in \Sigma$. If $\llbracket A \rrbracket(s) = true$, we say that A holds for s . When $\llbracket A \rrbracket(s) = true$ for all states $s \in \Sigma$, A is said to be *valid*, written $\models A$. For a set of assertions \mathcal{M} , we write $\models \mathcal{M}$ if $\models A$ holds for every $A \in \mathcal{M}$.

We have mentioned before that we will admit extensions of the assertion language with arbitrary functions and/or predicates that are not present in the programming language expressions. In this setting, the semantic interpretation of assertions must be complemented with a theory provided by the user, typically in the form of a set of axioms.

The semantic interpretation of a Hoare triple is a truth value that is independent of states. However there is a quantification over the set of states. Notice that a command C trivially satisfies a specification when the execution of C fails to terminate (since there exists no s such that $(C, s) \Downarrow s'$). A Hoare triple

Expressions:

$$\begin{aligned}
\llbracket n \rrbracket(s) &= n \quad \text{for } n \in \{\dots, -2, -1, 0, 1, 2, \dots\} \\
\llbracket x \rrbracket(s) &= s(x) \\
\llbracket -e \rrbracket(s) &= -\llbracket e \rrbracket(s) \\
\llbracket e_1 \square e_2 \rrbracket(s) &= \llbracket e_1 \rrbracket(s) \square \llbracket e_2 \rrbracket(s) \quad , \\
&\quad \text{for } (\square, \boxplus) \in \{(+, +), (-, -), (*, \times), (\mathbf{div}, \div), (\mathbf{mod}, \bmod)\} \\
\llbracket \mathbf{true} \rrbracket(s) &= \mathit{true} \\
\llbracket \mathbf{false} \rrbracket(s) &= \mathit{false} \\
\llbracket !e \rrbracket(s) &= \neg \llbracket e \rrbracket(s) \\
\llbracket e_1 \square e_2 \rrbracket(s) &= \llbracket e_1 \rrbracket(s) \square \llbracket e_2 \rrbracket(s) \quad , \\
&\quad \text{for } (\square, \boxplus) \in \{(\mathbf{==}, =), (\mathbf{!=}, \neq), (<, <), (<=, \leq), (>, >), (>=, \geq), (\&\&, \wedge), (\|\vee, \vee)\}
\end{aligned}$$

Commands:

1. $(\mathbf{skip}, s) \Downarrow s$
2. $(x := e, s) \Downarrow s[x \mapsto \llbracket e \rrbracket(s)]$
3. If $(C_1, s) \Downarrow s'$ and $(C_2, s') \Downarrow s''$ then $(C_1; C_2, s) \Downarrow s''$
4. If $(C_t, s) \Downarrow s'$ and $\llbracket b \rrbracket(s) = \mathit{true}$ then $(\mathbf{if } b \mathbf{ then } C_t \mathbf{ else } C_f, s) \Downarrow s'$
5. If $(C_f, s) \Downarrow s'$ and $\llbracket b \rrbracket(s) = \mathit{false}$ then $(\mathbf{if } b \mathbf{ then } C_t \mathbf{ else } C_f, s) \Downarrow s'$
6. If $\llbracket b \rrbracket(s) = \mathit{false}$ then $(\mathbf{while } b \mathbf{ do } C, s) \Downarrow s$
7. If $(C, s) \Downarrow s'$, $(\mathbf{while } b \mathbf{ do } C, s') \Downarrow s''$, and $\llbracket b \rrbracket(s) = \mathit{true}$ then $(\mathbf{while } b \mathbf{ do } C, s) \Downarrow s''$

Assertions:

$$\begin{aligned}
\llbracket \mathbf{true} \rrbracket(s) &= \mathit{true} \\
\llbracket \mathbf{false} \rrbracket(s) &= \mathit{false} \\
\llbracket !A \rrbracket(s) &= \neg \llbracket A \rrbracket(s) \\
\llbracket A_1 \square A_2 \rrbracket(s) &= \llbracket A_1 \rrbracket(s) \square \llbracket A_2 \rrbracket(s) \quad , \text{ for } (\square, \boxplus) \in \{(\&\&, \wedge), (\|\vee, \vee), (\rightarrow, \Rightarrow)\} \\
\llbracket \mathbf{forall } x. A \rrbracket(s) &= \forall v \in \mathbf{int}. \llbracket A \rrbracket(s[x \mapsto v]) \quad , \text{ with } v \text{ fresh} \\
\llbracket \mathbf{exists } x. A \rrbracket(s) &= \exists v \in \mathbf{int}. \llbracket A \rrbracket(s[x \mapsto v]) \quad , \text{ with } v \text{ fresh} \\
\llbracket e_1 \square e_2 \rrbracket(s) &= \llbracket e_1 \rrbracket(s) \square \llbracket e_2 \rrbracket(s) \quad , \\
&\quad \text{for } (\square, \boxplus) \in \{(\mathbf{==}, =), (\mathbf{!=}, \neq), (<, <), (<=, \leq), (>, >), (>=, \geq)\}
\end{aligned}$$

Specifications:

$$\llbracket \{P\} C \{Q\} \rrbracket = \forall s, s' \in \Sigma. \llbracket P \rrbracket(s) \wedge (C, s) \Downarrow s' \Rightarrow \llbracket Q \rrbracket(s')$$

Fig. 2. Semantic equations and natural semantics for While

$\{P\} C \{Q\}$ is valid if it is interpreted as true, i.e., for all states satisfying P , executing C either fails to terminate or terminates in a state satisfying Q .

Total Correctness. The above notion corresponds to *partial correctness* specifications, since termination is not guaranteed. If termination is required, we are in the presence of a *total correctness* formula and write $[P] C [Q]$ instead. The idea is that $[P] C [Q]$ is true if and only if, for all states satisfying

P , executing C terminates in a state satisfying Q . Note that the validity of a total correctness specification can be established by proving the corresponding partial correctness specification and additionally proving termination.

We will leave the discussion of termination out of this paper, since it is rather different in nature. Most program verification tools allow the users to specify *loop variants*: expressions whose value strictly decreases with each iteration, according to some well-founded relation. Typically non-negative integer expressions are used [42].

3 Hoare Logic

Given a specification in the form of a Hoare triple $\{P\} C \{Q\}$, how can its validity be checked? One could think of *testing* the program by running it with a battery of initial states satisfying the precondition P , and checking whether Q is satisfied after execution if the program terminates. Of course, this process cannot be exhaustive in general, and one can only expect to have a certain degree of confidence about the validity of a specification. This would not provide a proof of correctness.

The usual method for assuring the validity of specifications is to use a sound program-proof system. By sound we mean that it should not infer specifications that are not valid. We will now outline a formal system for inferring valid specifications and describe methods for mechanically verifying the correctness of given specifications.

A program-proof system is a set of *inference rules* that can be seen as fundamental laws about programs. Before defining an inference system for our While-language programs, let us first explain some basic concepts.

An *inference rule* consists of zero or more *premises* and a single *conclusion* (we use the notation of separating the premises from the conclusion by a horizontal line). Each rule, consisting in a number of premises and a conclusion, is in fact a *schema* for a specification (that is, a pattern containing meta-variables, each ranging over some phrase type). An *instance* of an inference rule is obtained by replacing all occurrences of each meta-variable by a phrase in its range. In some rules, there may be side conditions that must be satisfied by the replacement. Also, there may be syntactic operations (such as substitutions) that must be carried out after the replacement. An inference rule containing no premises is called an *axiom schema* (or simply, *axiom*). An inference rule is *sound* if and only if, for every instance, if the premises are all valid then the conclusion is valid.

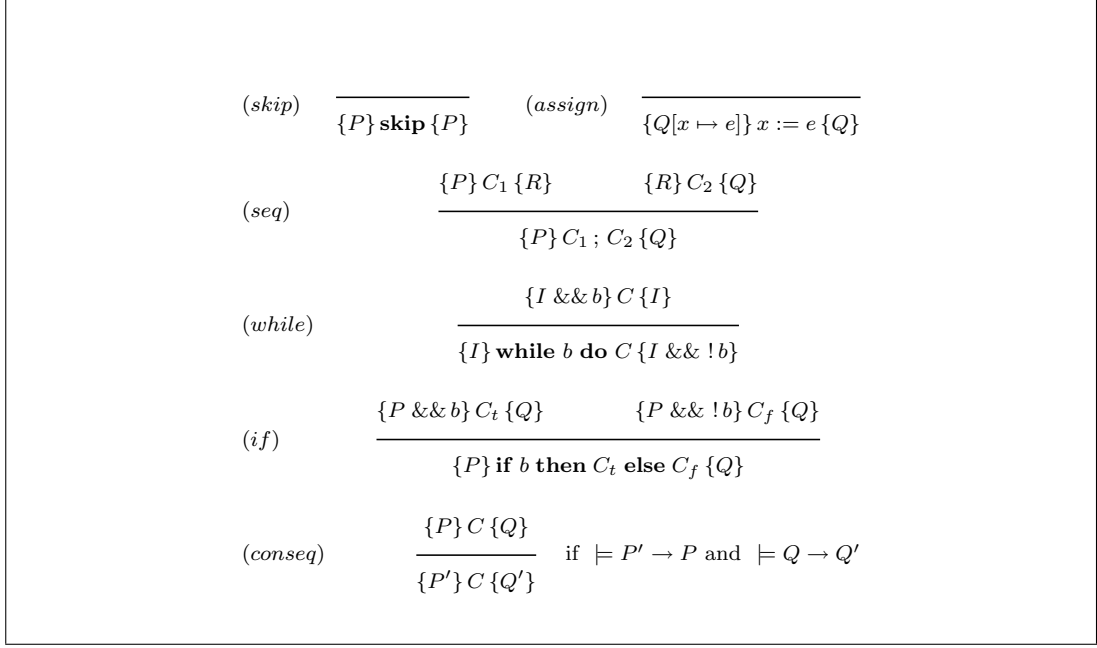


Fig. 3. Inference system of Hoare logic – System H

Deductions will be presented as trees; nodes will be labelled with formulas (specifications); the specification being proved lies at the root node, and the immediate subtrees are deductions corresponding to the premises of the rule instance used to infer the root specification formula. A *deduction tree* is also called a *derivation tree* or *proof tree*, and it constitutes a formal proof of the specification at its root.

The inference system of Hoare logic is shown in Figure 3. We name it system H. Moreover, we say that a Hoare triple $\{P\} C \{Q\}$ is *derivable* in this system, and write $\vdash_{\text{H}} \{P\} C \{Q\}$, if it can be inferred from the rules of Figure 3 (in general we will subscript the symbol \vdash with an inference system to denote derivability in that system).

This system comprises one rule for each command construct in the programming language, and the consequence rule, which allows us to derive a specification from another specification by strengthening the precondition or weakening the postcondition. Two of the command rules (*skip* and *assign*) are in fact axioms (since the do-nothing and assignment commands have no subcommands, the rules have no premises).

The assignment rule states that a postcondition Q can be granted for a command $x := e$ if the condition that results from substituting e for x in Q holds as precondition. $Q[x \mapsto e]$ stands for the result of substituting e for the free occurrences of x in Q .² The sequencing and conditional constructs are han-

² In the application of a substitution to a term, we rely on a variable convention.

dled by rules that are quite straightforward to understand. The rule for while commands uses the familiar notion of an *invariant* condition (denoted I in the rule), preserved by executions of the loop body. Note that one does not require that condition I holds *throughout* execution of the loop body; only that it is reestablished at the end of each iteration. We remark that, as the assertion language subsumes Boolean expressions, the Boolean condition b can be combined with assertions in the (*if*) and (*while*) rules (otherwise an embedding function would have to be applied to b).

Finally the (*conseq*) rule establishes a connection with predicate logic by means of side conditions that are assertions rather than specifications. The idea is that a specification can be derived from another specification provided that their corresponding preconditions and postconditions are related in the way dictated by the side conditions.

As a trivial example of using this system, consider the following proof tree concerning a single-instruction program, which is built from an instance of rule schema (*conseq*) and an instance of axiom schema (*assign*).

$$\frac{\overline{\{x + 1 > 10\} x := x + 1 \{x > 10\}}}{\{x > 10\} x := x + 1 \{x > 10\}} \text{ if } \models x > 10 \rightarrow x + 1 > 10 \text{ and } \models x > 10 \rightarrow x > 10$$

We remark that the side conditions concern the validity of purely first-order formulas (with no occurrences of program constructs). The compositional rules, associated with the program constructs of the language, depend on the semantics of commands but are *independent of the interpretation of data types and arithmetic operations*. It is the rule of consequence that brings data-specific assertions to bear on the proofs of specifications, through the introduction of side conditions.

Different presentations of the consequence rule can be found in the literature; some authors prefer to include the side conditions as premises in the rule. Observe however that care must be taken interpreting the following very common presentation of the rule

$$\frac{P' \rightarrow P \quad \{P\} C \{Q\} \quad Q \rightarrow Q'}{\{P'\} C \{Q'\}}$$

In particular, program variables that occur in $P' \rightarrow P$ and $Q \rightarrow Q'$ must be seen as universally quantified *locally* to those formulas, and not globally in the rule, which would be the usual interpretation in a natural deduction-like presentation.

The action of a substitution over a term is defined, as usual, with possible renaming of bound variables.

Because of the presence of the consequence rule, Hoare logic is not meant to be used by itself; it must always be accompanied by some device for establishing the validity of side conditions, such as a decision procedure based on satisfiability, or an inference system for first-order logic (usually known as predicate calculus). Naturally, this device should be adequate for reasoning about the concrete language of expressions of the programming language (which in our case implies reasoning with integer arithmetics), also taking into account, if applicable, user-provided axioms concerning additional functions and predicates of the assertion language.

Finally, note that it is an immediate consequence of the above discussion that reasoning about programs (with a first-order assertion language) is in general not decidable.

We will now address the relationship between system \mathbf{H} and the semantics of the language. All of the subsequent material relies on the fact that the inference system is sound: if some Hoare triple can be proved using system \mathbf{H} , then it is indeed valid according to the semantics.

Proposition 2 (Soundness) *In system \mathbf{H} every derivable specification is semantically valid. That is, if $\vdash_{\mathbf{H}} \{P\} C \{Q\}$, then $\llbracket \{P\} C \{Q\} \rrbracket = \text{true}$.*

Proof. By induction on the derivation of $\vdash_{\mathbf{H}} \{P\} C \{Q\}$. For the while case we also proceed by induction on the definition of the evaluation relation. \square

The classic reference on Hoare logic is [45]. The original ideas of Hoare had their roots in the work of Floyd, who had proposed a similar system for flowcharts [38]. For this reason Hoare logic is also known as Floyd-Hoare logic. Several papers survey its technical development [2, 21] and its historical development and impact [54]. In this paper we focus on applying Hoare logic in the context of program verification; its use in the software development process is covered in detail in textbooks [4, 81].

Bertot [11] has formalised in the *calculus of inductive constructions* the semantics (both operational and denotational) of a language like the one considered in this paper. A so-called *deep embedding* of the inference system of Hoare logic is given (the rules are encoded as cases of the inductive definition of a predicate), and correctness with respect to the operational semantics is proven. This very instructive work is available as part of an integrated Coq development (see Section 11) that covers many aspects of the semantics of programming languages. We will come back to this development in Section 7.

4 Verification Conditions and VCGens

With the help of a theorem prover or proof assistant, Hoare logic can be put into practice to produce a program verification system. This can be done in two ways. The first is by directly encoding the inference system in the logic of the proof tool and reasoning simultaneously with rules of both Hoare logic and first-order logic as required: reasoning starts with the former but switches to the latter logic when side conditions are introduced by the consequence rule.

The alternative approach, which is prevalent in modern program verification systems, is organised in two steps as follows:

- (1) A proof tree is constructed for the desired specification, *assuming that the side conditions generated by the consequence rule are valid.*
- (2) An external proof tool is used (such as a theorem prover or a proof assistant) to actually establish the validity of the side conditions.

We let $\mathcal{A} \Vdash_{\mathbf{H}} \{P\} C \{Q\}$, where \mathcal{A} represents a set of first-order assertions, denote the fact that there exists a derivation tree of $\{P\} C \{Q\}$ in system \mathbf{H} such that \mathcal{A} is the set of side conditions in that tree (in general we will subscript the symbol \Vdash with an inference system to denote the existence of a proof tree in that system with a given set of side conditions).

Note that the soundness of this approach to verification is immediate: if $\mathcal{A} \Vdash_{\mathbf{H}} \{P\} C \{Q\}$ and all the assertions of \mathcal{A} hold, then $\vdash_{\mathbf{H}} \{P\} C \{Q\}$. The first step is said to generate *proof obligations* (the elements of \mathcal{A}) that must then be *discharged* in the second step. In the context of program verification these proof obligations are usually known as *verification conditions*.

The advantage of this method lies on its flexibility. Since program constructs do not occur in the assertions in \mathcal{A} , the second step involves only discharging first-order proof obligations, for which a great number of proof tools can be used interchangeably or even cooperatively. It is also easier to modify a program verification system organized in this way. For instance if the programming language is modified, only the first step above is affected.

Clearly the inference system allows for different proof trees to be constructed for the same conclusion specification, and in fact these trees may well have different sets of side conditions. This tree construction process can be replaced by a simple algorithm – a *Verification Conditions Generator* (VCGen) – that constructs a set of verification conditions by applying a specific strategy, i.e. the algorithm produces verification conditions that correspond to the side conditions of one particular derivation.

VCGen algorithms will be given as functions that take as input a Hoare triple

and return a set of first-order proof obligations. A side condition of the form $\models P \rightarrow Q$ will give rise to a verification condition written as $[P \rightarrow Q]$, where the $[\cdot]$ notation represents the processing that may be required to export proof obligations to the target proof tool.

This processing may require more than just translating assertions into the language of the tool; one typical operation corresponds to calculating the *universal closure* of an assertion, by making explicit the universal quantification over program variables that is implicit in the notion of validity (i.e. truth in all states) of side conditions. For instance, the side condition $\models x > 10 \rightarrow x + 1 > 10$ in the earlier example would generate the verification condition $[x > 10 \rightarrow x + 1 > 10]$, which in turn could be exported as the formula `forall x. x > 10 -> x + 1 > 10`.

We remark that there are two possible sources of errors that may cause the verification of a given Hoare triple to fail. These are:

- (1) *program errors*; and
- (2) *specification errors*: the program may be correct with respect to the intended specification, but the specification has not been correctly formalised (there are errors in the preconditions or postconditions).

In the rest of this paper we will study several VCGen algorithms and show how they are obtained from the inference system \mathbf{H} or some other related system. Our first step is to study how the construction of derivations of \mathbf{H} can be mechanised.

5 An Alternative Formulation of Hoare Logic

We now focus on using Hoare logic to produce verification conditions, as part of the verification architecture outlined in the previous section. Given a specification, we wish to construct a proof tree having that specification as conclusion; verification conditions result from the side conditions of instances of rules in that tree.

There are two desirable properties that the inference system of Hoare logic should enjoy to make possible the automatic construction of proof trees.

- (1) The *subformula* property: the premises of a rule should not contain occurrences of assertions that do not occur in the rule's conclusion. In other words, all the assertions that occur in the premises should be subformulas of those occurring in the conclusion. Otherwise one would have to invent formulas when applying the rule in a backward fashion.

$$\begin{array}{c}
\frac{}{\{P\} \mathbf{skip} \{Q\}} \quad \text{if } \models P \rightarrow Q \qquad \frac{}{\{P\} x := e \{Q\}} \quad \text{if } \models P \rightarrow Q[x \mapsto e] \\
\\
\frac{\{P\} C_1 \{R\} \qquad \{R\} C_2 \{Q\}}{\{P\} C_1 ; C_2 \{Q\}} \\
\\
\frac{\{I \ \&\& \ b\} C \{I\}}{\{P\} \mathbf{while} \ b \ \mathbf{do} \ C \{Q\}} \quad \text{if } \models P \rightarrow I \ \text{and} \ \models I \ \&\& \ !b \rightarrow Q \\
\\
\frac{\{P \ \&\& \ b\} C_t \{Q\} \qquad \{P \ \&\& \ !b\} C_f \{Q\}}{\{P\} \mathbf{if} \ b \ \mathbf{then} \ C_t \ \mathbf{else} \ C_f \{Q\}}
\end{array}$$

Fig. 4. Goal-directed version of Hoare logic – System \mathbf{Hg}

- (2) Unambiguity: a unique rule should be applicable in a backward fashion for any given goal, so that the construction of derivation trees can be syntax-directed.

We start with the second property. The system \mathbf{H} of Hoare logic can easily be transformed into an equivalent unambiguous system; observe that ambiguity is only caused by the presence of the consequence rule. In particular, note that some rules are only applicable to goals that satisfy certain constraints, namely:

- the (*skip*) rule can only be applied if the precondition and the postcondition are equal;
- the (*assign*) rule for assignment can only be applied if the precondition results from the postcondition by performing the corresponding substitution;
- the (*while*) rule can only be applied if the precondition is an invariant of the loop, and the postcondition is the same invariant strengthened with the negation of the loop condition.

Application of these three rules to goals with arbitrary preconditions and postconditions may previously require the application of the consequence rule. This becomes unnecessary if the weakening/strengthening conditions are judiciously distributed through the program rules. This is precisely the case in the system of Figure 4.

This new system may be called a *goal-directed* system, since it consists of exactly one rule for each program construct, and moreover, for a given specification $\{P\} C \{Q\}$, the rule matching the program C can *always* be applied (if the side conditions are met), since the precondition and postcondition are now arbitrary. We name it system \mathbf{Hg} , and we let $\vdash_{\mathbf{Hg}} \{P\} C \{Q\}$ denote the

fact that $\{P\} C \{Q\}$ can be inferred from the rules of Figure 4.

It is easy to see that the (*conseq*) rule is admissible in system **Hg**, and that systems **H** and **Hg** are equivalent.

Lemma 3 *If $\vdash_{\mathbf{Hg}} \{P\} C \{Q\}$ and both $\models P' \rightarrow P$ and $\models Q \rightarrow Q'$ hold, then $\vdash_{\mathbf{Hg}} \{P'\} C \{Q'\}$.*

Proof. By induction on the derivation of $\vdash_{\mathbf{Hg}} \{P\} C \{Q\}$. □

Proposition 4 $\vdash_{\mathbf{H}} \{P\} C \{Q\}$ *iff* $\vdash_{\mathbf{Hg}} \{P\} C \{Q\}$.

Proof. \Rightarrow) By induction on the derivation of $\vdash_{\mathbf{H}} \{P\} C \{Q\}$, using Lemma 3.
 \Leftarrow) By induction on the derivation of $\vdash_{\mathbf{Hg}} \{P\} C \{Q\}$. □

We remark that this new system still does not enjoy the subformula property. In fact, although we have removed two rules that did not enjoy the property, we have also caused the while rule to lose that property since the invariant assertion no longer occurs in the conclusion. Moreover the rule for the sequence construct does not enjoy the property either, since the intermediate assertion R does not occur in the conclusion.

In fact, attaining the subformula property, and the possibility of automated proof construction, requires the programmer to provide some extra information.

6 Program Annotations

A fully automated verification process would require mechanisation at different levels. In particular, one would have to be able to generate loop invariants and to establish the validity of first-order conditions mechanically (admittedly, an impossible goal). In Section 11 we discuss advances in both of these areas. For now, we concentrate on another, much easier aspect, which is automating the process of reasoning with Hoare logic.

One way to restore the subformula property in our current system for Hoare logic is precisely to introduce human-provided *annotations* in the programs. An *annotated program* is a program with assertions embedded within it. Inserted assertions should express conditions one expects to hold whenever control reaches the points at which they occur. Let **AComm** be the class of

$$\begin{array}{c}
\frac{}{\{P\} \mathbf{skip} \{Q\}} \quad \text{if } \models P \rightarrow Q \qquad \frac{}{\{P\} x := e \{Q\}} \quad \text{if } \models P \rightarrow Q[x \mapsto e] \\
\\
\frac{\{P\} C_1 \{R\} \qquad \{R\} C_2 \{Q\}}{\{P\} C_1 ; \{R\} C_2 \{Q\}} \\
\\
\frac{\{I \ \&\& \ b\} C \{I\}}{\{P\} \mathbf{while} \ b \ \mathbf{do} \ \{I\} C \{Q\}} \quad \text{if } \models P \rightarrow I \ \text{and} \ \models I \ \&\& \ !b \rightarrow Q \\
\\
\frac{\{P \ \&\& \ b\} C_t \{Q\} \qquad \{P \ \&\& \ !b\} C_f \{Q\}}{\{P\} \mathbf{if} \ b \ \mathbf{then} \ C_t \ \mathbf{else} \ C_f \{Q\}}
\end{array}$$

Fig. 5. Goal-directed version of Hoare logic for annotated programs – System **Hga**. *annotated commands*. Its abstract syntax is defined by

AComm $\ni C ::= \mathbf{skip} \mid C ; \{A\} C \mid x := e \mid \mathbf{if} \ b \ \mathbf{then} \ C \ \mathbf{else} \ C \mid \mathbf{while} \ b \ \mathbf{do} \ \{A\} C$

For instance the annotated command

while b **do** $\{I\} C$

denotes a loop with condition b , instruction body C , and (user-provided) invariant I . Naturally, the introduction of the loop invariant plays no role in the execution semantics of the loop. Similarly, the command sequence

$C_1 ; \{R\} C_2$

has the same semantics as $C_1 ; C_2$, but is annotated with an assertion that must be true when the execution of C_1 (started in a state in which any given preconditions of the combined command hold) terminates.

An inference system of Hoare logic for annotated programs is shown in Figure 5. We name it system **Hga**. Note that this system enjoys the subformula property, and can be used mechanically in a backward fashion to generate verification conditions. We will exemplify its use with a classic program. This example also illustrates the need to extend the assertion language with a vocabulary of function symbols.

Example 5 Consider the program for calculating Fibonacci numbers in Figure 6. The annotated program C_{Fib}^A shown in b) is the result of annotating C_{Fib} shown in a) in an ad hoc way.

a) Program C_{Fib}	b) Annotated program C_{Fib}^A
<pre> x := 1; y := 0; i := 1; while i < n do { aux := y; y := x; x := x + aux; i := i + 1 } </pre>	<pre> x := 1; {x == 1} y := 0; {x == 1 && y == 0} i := 1; {x == 1 && y == 0 && i == 1} while i < n do {i ≤ n && x == Fib(i) && y == Fib(i - 1)} { aux := y; {i ≤ n && x == Fib(i) && aux == Fib(i - 1)} y := x; {i < n && x == Fib(i) && y == Fib(i) && aux == Fib(i - 1)} x := x + aux; {i ≤ n && x == Fib(i) + Fib(i - 1) && y == Fib(i)} i := i + 1; } </pre>

Fig. 6. Example program: Fibonacci

We extend the vocabulary of the assertion language with a function `Fib` with arity 1. This function will be transparent from the point of view of verification condition generation, and it may naturally occur in the proof obligations. To give meaning to this function, we provide a theory consisting of the following axioms:

$$\text{Fib}(0) == 0$$

$$\text{Fib}(1) == 1$$

$$\text{Forall } x. x > 1 \rightarrow \text{Fib}(x) == \text{Fib}(x - 1) + \text{Fib}(x - 2)$$

Observe that we could have chosen to extend the language with a predicate of arity 2 instead, where the second argument stands for the Fibonacci number of the first. We also remark the following:

- The invariant states that variables x and y are used to store respectively the Fibonacci numbers for i and $i - 1$, together with an obvious condition regarding a bound for i .
- The annotations in the sequence of assignment instructions preceding the loop basically store the current state of the program.
- Inside the loop body, the annotated state is reset, since all that is known when execution enters this sequence of instructions in a particular iteration

is that the invariant was initially valid. The annotations are then propagated forward from the invariant with each assignment instruction.

- The annotations are somewhat optimised. For instance after the first instruction in the loop body, since y is about to lose its present value, there is no need to keep the old value in the annotation. Similarly, aux is dropped from the annotation as soon as it is clear that it will no longer be required.

Let us consider the following specification for this annotated program C_{Fib}^A

$$\{n > 0\} C_{\text{Fib}}^A \{x == \text{Fib}(n)\}$$

Suppose we take composition of programs to be right-associative and let C be the body of the loop; then applying the rule for the sequence construct we obtain the following two proof obligations

$$\begin{aligned} &\{n > 0\} \\ &x := 1; \{x == 1\} y := 0; \{x == 1 \ \&\& \ y == 0\} i := 1 \\ &\{x == 1 \ \&\& \ y == 0 \ \&\& \ i == 1\} \end{aligned}$$

and

$$\begin{aligned} &\{x == 1 \ \&\& \ y == 0 \ \&\& \ i == 1\} \\ &\mathbf{while} \ i < n \ \mathbf{do} \ \{i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i - 1)\} C \\ &\{x == \text{Fib}(n)\} \end{aligned}$$

The former will generate trivial verification conditions, and for the latter, applying the loop rule yields two verification conditions, namely

$$\models x == 1 \ \&\& \ y == 0 \ \&\& \ i == 1 \rightarrow i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i - 1)$$

and

$$\models i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i - 1) \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)$$

together with the proof obligation corresponding to the loop invariant preservation,

$$\{i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i - 1) \ \&\& \ i < n\} C \{i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i - 1)\}$$

The reader is invited to continue this example through to the end. We will return to it in subsequent sections.

Note that if the annotations introduced in a program are ‘wrong’ (for instance, the user provides as loop invariant a property that is not preserved by the loop body), there is no risk of corrupting the soundness of the verification process: simply, proof obligations will be created that cannot be proved.

Each choice of annotations leads to the construction of a unique tree in system Hga for a given Hoare triple (possibly a proof tree if the side conditions are

valid). In system **Hg** for non-annotated programs, many different proof trees are admissible for the same specification. However, if a Hoare triple for an annotated program is derivable in system **Hga**, then the corresponding triple for the program without annotations is derivable in system **Hg**. To formally state this property we first define an erasure function $\text{er}_A : \mathbf{AComm} \rightarrow \mathbf{Comm}$ inductively as follows

$$\begin{aligned} \text{er}_A(\text{skip}) &= \text{skip} \\ \text{er}_A(x := e) &= x := e \\ \text{er}_A(C_1; \{R\} C_2) &= \text{er}_A(C_1); \text{er}_A(C_2) \\ \text{er}_A(\text{if } b \text{ then } C_t \text{ else } C_f) &= \text{if } b \text{ then } \text{er}_A(C_t) \text{ else } \text{er}_A(C_f) \\ \text{er}_A(\text{while } b \text{ do } \{I\} C) &= \text{while } b \text{ do } \text{er}_A(C) \end{aligned}$$

Proposition 6 *If $\vdash_{\text{Hga}} \{P\} C \{Q\}$, then $\vdash_{\text{Hg}} \{P\} \text{er}_A(C) \{Q\}$.*

Proof. By induction on the derivation of $\vdash_{\text{Hga}} \{P\} C \{Q\}$. □

We say that a program C is *correctly annotated with respect to a specification* (P, Q) if $\vdash_{\text{Hga}} \{P\} C \{Q\}$ whenever $\vdash_{\text{Hg}} \{P\} \text{er}_A(C) \{Q\}$.

Note that additionally to the sources of errors identified in Section 4, annotations provide a new opportunity for errors to occur. *Annotation errors* occur when a program has been annotated with assertions that may not hold when the corresponding program point is reached during execution, which may well cause the verification of the program to fail.

It is clear from the above example that annotating programs is a mostly tedious activity. This is true in particular for intermediate assertions in sequences of commands. Fortunately, these assertions can easily be inferred mechanically, if the program is seen in the context of a specification. Consider for instance the Hoare triple

$$\begin{aligned} &\{x == 5 \ \&\& \ y == 10\} \\ &\quad aux := y; \\ &\quad y := x; \\ &\quad x := x + aux; \\ &\{x > 10 \ \&\& \ y == 5\} \end{aligned}$$

One possible correct way to annotate this sequence of commands is to work backwards from the postcondition to obtain a precondition for each atomic command (in this example it suffices to use the assignment rule). The alternative would be to propagate the precondition forward. Figure 7, a) and b), illustrates these two approaches.

The two annotation strategies exemplified in Figure 7 will in the next two sections be explored and give rise to two different VCGen algorithms, that




Fig. 7. Annotation propagation




Fig. 8. Goal-directed Hoare logic with annotated while-loops – System **Hgi**.

work on programs annotated with loop invariants only. We shall now state some lemmas concerning such programs.

We let **IComm** denote the class of programs with annotated while commands. Its abstract syntax is defined by

$$\mathbf{IComm} \ni C ::= \mathbf{skip} \mid C ; C \mid x := e \mid \mathbf{if} \ b \ \mathbf{then} \ C \ \mathbf{else} \ C \mid \mathbf{while} \ b \ \mathbf{do} \ \{A\} \ C$$

An inference system of Hoare logic for **IComm** programs is shown in Figure 8. We call it system **Hgi**.

Of course, every derivation of a Hoare triple in **Hgi** has a correspondence in **Hg**. To state this property formally we define an erasure function $\text{er}_1 : \mathbf{IComm} \rightarrow$

Comm similar to er_A (with the obvious adaptation).

Proposition 7 *If $\vdash_{\text{Hgi}} \{P\} C \{Q\}$, then $\vdash_{\text{Hg}} \{P\} \text{er}_I(C) \{Q\}$.*

Proof. By induction on the derivation of $\vdash_{\text{Hgi}} \{P\} C \{Q\}$. \square

For $C \in \mathbf{IComm}$, $P, Q \in \mathbf{Assert}$, we say that C is *correctly annotated wrt.* (P, Q) if $\vdash_{\text{Hgi}} \{P\} C \{Q\}$ whenever $\vdash_{\text{Hg}} \{P\} \text{er}_I(C) \{Q\}$.

Lemma 8 *If $\vdash_{\text{Hgi}} \{P\} C \{Q\}$, then C is correctly annotated wrt. (P, Q) .*

Proof. Trivial, regarding the definition of correctly annotated program. \square

Lemma 9 *If C is correctly annotated wrt. (P, Q) and both $\models P' \rightarrow P$ and $\models Q \rightarrow Q'$ hold, then C is correctly annotated wrt. (P', Q') .*

Proof. Immediate from Lemma 3. \square

Lemma 10

- (1) *If $\vdash_{\text{Hg}} \{P\} \text{er}_I(C_1); \text{er}_I(C_2) \{Q\}$ and $C_1; C_2$ is correctly annotated wrt. (P, Q) , then, for some assertion R , C_1 is correctly annotated wrt. (P, R) , C_2 is correctly annotated wrt. (R, Q) and both $\vdash_{\text{Hg}} \{P\} \text{er}_I(C_1) \{R\}$ and $\vdash_{\text{Hg}} \{R\} \text{er}_I(C_2) \{Q\}$ hold.*
- (2) *If $\vdash_{\text{Hg}} \{P\} \mathbf{if} \ b \ \mathbf{then} \ \text{er}_I(C_t) \ \mathbf{else} \ \text{er}_I(C_f) \ \{Q\}$ and $\mathbf{if} \ b \ \mathbf{then} \ C_t \ \mathbf{else} \ C_f$ is correctly annotated wrt. (P, Q) , then C_t is correctly annotated wrt. $(P \ \&\& \ b, Q)$, C_f is correctly annotated wrt. $(P \ \&\& \ !b, Q)$ and both $\vdash_{\text{Hg}} \{P \ \&\& \ b\} \text{er}_I(C_t) \ \{Q\}$ and $\vdash_{\text{Hg}} \{P \ \&\& \ !b\} \text{er}_I(C_f) \ \{Q\}$ hold.*

Proof. Both (1) and (2) are proved by definition of correctly annotated program, case analysis on the rules of system **Hgi**, Lemma 8 and Proposition 7. \square

Gordon [41] has proposed a mechanisation of Hoare logic in the HOL proof assistant, which includes derivations of the inference rules of Hoare logic from a semantic description of the language. A system that is close to system **Hga** is also proposed.

In Gordon's system these rules are not seen as inference rules; they are instead used to define *tactics* for the prover. A tactic is a function used to advance the proof construction, i.e. it is applied to the current proof state to produce subgoals of the present goal. In Gordon's system the VCGen is itself implemented as a tactic. The system incorporates notions that will be explained in

$$\begin{aligned}
\text{wprec}(\text{skip}, Q) &= Q \\
\text{wprec}(x := e, Q) &= Q[x \mapsto e] \\
\text{wprec}(C_1; C_2, Q) &= \text{wprec}(C_1, \text{wprec}(C_2, Q)) \\
\text{wprec}(\text{if } b \text{ then } C_t \text{ else } C_f, Q) &= (b \rightarrow \text{wprec}(C_t, Q)) \ \&\& \ (!b \rightarrow \text{wprec}(C_f, Q)) \\
\text{wprec}(\text{while } b \text{ do } \{I\} C, Q) &= I
\end{aligned}$$

Fig. 9. Definition of weakest precondition for **IComm** programs: wprec

the next section, so we will return to it.

Gordon uses a so-called *shallow embedding* of the language into the proof system’s logic, which precludes proving the assignment axiom as a HOL theorem (it is a meta-level property). Homeier and Martin [47] use a *deep embedding* instead and achieve a proof of correctness for a VCGen; their work very likely reports the first *fully verified VCGen*.

The difference between a shallow and a deep embedding is dictated by the way in which the represented languages (for programs and assertions) are related to the object language of the theorem prover – either as extensions to the latter, or constructed with separate data types. See [3] for a survey on representing Hoare logic in the language of a theorem prover, and the different embedding possibilities.

7 Weakest Preconditions

Backward propagation of assertions can be realized through the use of *weakest preconditions*. Given a program $C \in \mathbf{IComm}$ and a postcondition Q , we calculate an assertion $\text{wprec}(C, Q)$ such that $\{\text{wprec}(C, Q)\} C \{Q\}$ is valid and moreover if $\{P\} C \{Q\}$ is valid for some P then $\models P \rightarrow \text{wprec}(C, Q)$. Thus $\text{wprec}(C, Q)$ is the weakest precondition that grants the truth of postcondition Q after terminating executions of C .

Weakest preconditions are implicit in Hoare logic, in the way that the assignment axiom propagates postconditions backwards: $Q[x \mapsto e]$ is the weakest precondition for Q to hold after execution of $x := e$. The explicit definition of wprec given in Figure 9 propagates postconditions backwards through the remaining program constructs.

In the sequence command, the weakest precondition of the second command is fed as postcondition to the first command, to obtain the precondition of

the combined command. The clause for conditional is also straightforward to understand: the weakest precondition of such a command is the weakest precondition of the appropriate branch (calculated considering the same post-condition), depending on the value of the Boolean condition.

In general, since the number of iterations may not be known at compile time, the weakest precondition of a loop cannot be calculated statically – performing a one-step expansion of a *while* command (using conditional) and trying to derive the weakest precondition from that expansion leads to a recursive equation. However, for the case of loops annotated with invariants, the weakest precondition can be defined to be precisely the invariant assertion. The reason for this is that all the reasoning about the behaviour of loops depends on the invariants being granted as preconditions. In an annotated program, the invariant of a loop is thus the weakest precondition required for any post-condition to hold.

Proposition 11 *If $C \in \mathbf{IComm}$ is correctly annotated wrt. (P, Q) and $\vdash_{\mathbf{Hgi}} \{P\} \text{er}_1(C) \{Q\}$, then*

- (1) $\vdash_{\mathbf{Hgi}} \{\text{wprec}(C, Q)\} C \{Q\}$
- (2) $\models P \rightarrow \text{wprec}(C, Q)$

Proof. By induction on the structure of C , and using lemmas 3, 8 and 10 for (1) and 3, 9 and 10 for (2). □

Figure 10 contains the straightforward definition of a VCGen obtained from system \mathbf{Hgi} by using wprec as an auxiliary function. The following lemma states that the verification conditions produced by this algorithm, for a given Hoare triple, are exactly the side conditions of a derivation tree of that triple in system \mathbf{Hgi} .

Lemma 12 *For any $C \in \mathbf{IComm}$, $\text{VC}(\{P\} C \{Q\}) \Vdash_{\mathbf{Hgi}} \{P\} C \{Q\}$*

Proof. By induction on the structure of C . □

Before looking at an example, we remark the following. The VCGen is implicitly constructing a derivation of \mathbf{Hg} following a fixed strategy. This consists in always constructing the subderivation corresponding to the second subcommand in any sequence command, until the intermediate condition is attained, at which point the first subderivation can then be constructed. This results in the derivation that would be constructed in \mathbf{Hga} if the program was previously annotated with the intermediate conditions calculated by wprec .

Example 13 *Figure 12 shows the application of the VCGen algorithm to our*

$$\begin{aligned}
\text{VC}(\{P\} \text{skip} \{Q\}) &= \{[P \rightarrow Q]\} \\
\text{VC}(\{P\} x := e \{Q\}) &= \{[P \rightarrow Q[x \mapsto e]]\} \\
\text{VC}(\{P\} C_1 ; C_2 \{Q\}) &= \text{VC}(\{P\} C_1 \{\text{wprec}(C_2, Q)\}) \cup \text{VC}(\{\text{wprec}(C_2, Q)\} C_2 \{Q\}) \\
\text{VC}(\{P\} \text{while } b \text{ do } \{I\} C \{Q\}) &= \{[P \rightarrow I], [I \ \&\& \ !b \rightarrow Q]\} \cup \text{VC}(\{I \ \&\& \ b\} C \{I\}) \\
\text{VC}(\{P\} \text{if } b \text{ then } C_t \text{ else } C_f \{Q\}) &= \text{VC}(\{P \ \&\& \ b\} C_t \{Q\}) \cup \text{VC}(\{P \ \&\& \ !b\} C_f \{Q\})
\end{aligned}$$

Fig. 10. A VCGen for **IComm** based on weakest preconditions: VC

$$\begin{array}{l}
C' \left\{ \begin{array}{l}
A_0 = \{1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)\} \\
x := 1; \\
A_1 = \{1 \leq n \ \&\& \ x == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)\} \\
y := 0; \\
A_2 = \{1 \leq n \ \&\& \ x == \text{Fib}(1) \ \&\& \ y == \text{Fib}(1-1)\} \\
i := 1; \\
I = \{i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1)\} \\
\text{while } i < n \text{ do } \{ i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1)\} \\
\left\{ \begin{array}{l}
\{ \\
aux := y; \\
A_3 = \{i+1 \leq n \ \&\& \ x + aux == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)\} \\
y := x; \\
A_4 = \{i+1 \leq n \ \&\& \ x + aux == \text{Fib}(i+1) \ \&\& \ y == \text{Fib}((i+1)-1)\} \\
x := x + aux; \\
A_5 = \{i+1 \leq n \ \&\& \ x == \text{Fib}(i+1) \ \&\& \ y == \text{Fib}((i+1)-1)\} \\
i := i + 1; \\
\} \\
\}
\end{array} \right. \\
C'' \left\{ \begin{array}{l}
\{ \\
aux := y; \\
A_3 = \{i+1 \leq n \ \&\& \ x + aux == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)\} \\
y := x; \\
A_4 = \{i+1 \leq n \ \&\& \ x + aux == \text{Fib}(i+1) \ \&\& \ y == \text{Fib}((i+1)-1)\} \\
x := x + aux; \\
A_5 = \{i+1 \leq n \ \&\& \ x == \text{Fib}(i+1) \ \&\& \ y == \text{Fib}((i+1)-1)\} \\
i := i + 1; \\
\} \\
\}
\end{array} \right.
\end{array}$$

Fig. 11. Fibonacci with annotated loop: $C_{\text{Fib}}^{\text{I}}$

running example program. The intermediate conditions (generated as weakest preconditions by `wprec`) are shown in Figure 11 as annotations in the program, for the sake of readability.

Observe that while all the conditions generated are straightforward to prove, most are tautological, of the form $[Q \rightarrow Q]$ for some Q . The ones that are not are

- (1) $[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)]$, which states that the precondition in the specification is stronger than the weakest precondition calculated for the program;
- (2) $[I \ \&\& \ i < n \rightarrow i + 1 \leq n \ \&\& \ x + y == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)]$

$$\begin{aligned}
& \text{VC}(\{n > 0\} C_{\text{Fib}}^I \{x == \text{Fib}(n)\}) \\
&= \text{VC}(\{n > 0\} x := 1 \{A_1\}) \cup \text{VC}(\{A_1\} y := 0 \{A_2\}) \cup \text{VC}(\{A_2\} i := 1 \{I\}) \cup \\
&\quad \text{VC}(\{I\} \text{while } i < n \text{ do } \{I\} C'' \{x == \text{Fib}(n)\}) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1 - 1)]\} \cup \\
&\quad \{[A_1 \rightarrow A_1]\} \cup \{[A_2 \rightarrow A_2]\} \cup \\
&\quad \{[I \rightarrow I], [I \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)]\} \cup \text{VC}(\{I \ \&\& \ i < n\} C'' \{I\}) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1 - 1)], \\
&\quad [A_1 \rightarrow A_1], [A_2 \rightarrow A_2], [I \rightarrow I], [I \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)], \\
&\quad [I \ \&\& \ i < n \rightarrow i + 1 \leq n \ \&\& \ x + y == \text{Fib}(i + 1) \ \&\& \ x == \text{Fib}((i + 1) - 1)], \\
&\quad [A_3 \rightarrow A_3], [A_4 \rightarrow A_4], [A_5 \rightarrow A_5]\}
\end{aligned}$$

Fig. 12. Fibonacci example: verification conditions obtained with VC.

1) $- 1]$, which corresponds to the preservation of the loop invariant; and
(3) $[I \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)]$, corresponding to the “use case” upon termination of the loop, which must imply the postcondition in the Hoare triple.

The reason why this VCGen generates many tautological conditions is that for any command of the form $C_1; x := e$ with postcondition Q , a precondition $\text{wprec}(x := e, Q)$ will be generated for the assignment command. Conditions like the following will proliferate

$$\text{VC}(\{\text{wprec}(x := e, Q)\} x := e \{Q\}) = [Q[x \mapsto e] \rightarrow Q[x \mapsto e]]$$

It is however possible to define a VCGen that eliminates these unnecessary conditions. The trick is to calculate verification conditions independently of preconditions. The VCGen in Figure 13, that we call VCG, uses an auxiliary recursive function VC_{aux} that does exactly this.

Note that for any sequence of assignments VCG generates a sole verification condition. This is a consequence of the following lemma.

Lemma 14 $\text{VC}_{\text{aux}}(C_1; C_2; \dots; C_n, Q) = \emptyset$, if each C_i is either a skip or an assignment.

Proof. By induction on the length of the sequence $C_1; C_2; \dots; C_n$. □

The verification conditions for a Hoare triple $\{P\} C \{Q\}$ are then calculated by adding to the set $\text{VC}_{\text{aux}}(C, Q)$ a condition explicitly relating P and the weakest precondition $\text{wprec}(C, Q)$. This may be seen as the principal verification condition, obtained from the backward propagation of the postcondition, whereas $\text{VC}_{\text{aux}}(C, Q)$ calculates secondary verification conditions resulting from loops.

$$\begin{aligned}
\text{VC}_{\text{aux}}(\text{skip}, Q) &= \emptyset \\
\text{VC}_{\text{aux}}(x := e, Q) &= \emptyset \\
\text{VC}_{\text{aux}}(C_1; C_2, Q) &= \text{VC}_{\text{aux}}(C_1, \text{wprec}(C_2, Q)) \cup \text{VC}_{\text{aux}}(C_2, Q) \\
\text{VC}_{\text{aux}}(\text{if } b \text{ then } C_t \text{ else } C_f, Q) &= \text{VC}_{\text{aux}}(C_t, Q) \cup \text{VC}_{\text{aux}}(C_f, Q) \\
\text{VC}_{\text{aux}}(\text{while } b \text{ do } \{I\} C, Q) &= \{[(I \ \&\& \ b) \rightarrow \text{wprec}(C, I)]\} \cup \text{VC}_{\text{aux}}(C, I) \cup \{[(I \ \&\& \ !b) \rightarrow Q]\} \\
\text{VCG}(\{P\} C \{Q\}) &= \{[P \rightarrow \text{wprec}(C, Q)]\} \cup \text{VC}_{\text{aux}}(C, Q)
\end{aligned}$$

Fig. 13. An improved VCGen: VCG

$$\begin{aligned}
&\text{VCG}(\{n > 0\} C_{\text{Fib}}^I \{x == \text{Fib}(n)\}) \\
&= \{[n > 0 \rightarrow \text{wprec}(C_{\text{Fib}}^I, x == \text{Fib}(n))]\} \cup \text{VC}_{\text{aux}}(C_{\text{Fib}}^I, x == \text{Fib}(n)) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)]\} \cup \text{VC}_{\text{aux}}(C', I) \cup \\
&\quad \text{VC}_{\text{aux}}(\text{while } i < n \text{ do } \{I\} C'', x == \text{Fib}(n)) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)]\} \cup \emptyset \cup \\
&\quad \{[I \ \&\& \ i < n \rightarrow \text{wprec}(C'', I)]\} \cup \text{VC}_{\text{aux}}(C'', I) \cup \{[I \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)]\} \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)], \\
&\quad [I \ \&\& \ i < n \rightarrow i + 1 \leq n \ \&\& \ x + y == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)], \\
&\quad [I \ \&\& \ !(i < n) \rightarrow x == \text{Fib}(n)]\}
\end{aligned}$$

Fig. 14. Fibonacci example, using VCG

It is easy to see that, for a given Hoare triple, the verification conditions generated by VCG entail the ones generated by VC.

Lemma 15 *If $\models \text{VCG}(\{P\} C \{Q\})$, then $\models \text{VC}(\{P\} C \{Q\})$.*

Proof. By induction on the structure of C . □

Proposition 16 (Correctness of VCG) *Let $C \in \text{IComm}$ and $P, Q \in \text{Assert}$ such that $\models \text{VCG}(\{P\} C \{Q\})$. Then $\vdash_{\text{Hg}} \{P\} \text{er}_I(C) \{Q\}$.*

Proof. Immediate by lemmas 12 and 15, and Proposition 7. □

Figure 14 illustrates the use of this VCGen. We remark that, since wprec and VC_{aux} perform similar traversals of the program structure, they could be fused into a single function using the well-known *tupling* technique of functional programming.

Bertot’s Coq development on program semantics [11] (see Section 3) includes both the definition of the VC as a recursive function, and its proof of correctness.

Gordon’s VCGen [41] (see Section 6) is very similar to the VCG presented in this section; the difference is that his algorithm does not require invoking an external function for calculating weakest preconditions; instead sequence commands are required to be partially annotated. A sequence command of the form $C_1 ; x := e$ does not need an annotation, but all other sequences have to be annotated with an intermediate assertion, as in $C_1 ; \{R\} C_2$. This system thus stands between Hga and Hgi.

Weakest preconditions were introduced by Dijkstra [29] from a semantic perspective. The idea was to interpret programs as *predicate transformers* – mappings of postconditions into preconditions. Dijkstra used a more abstract programming language based on *guarded commands*, and in fact some modern and very advanced tools for program verification are based on such a language. This will be the topic of Section 8.

A Note on Auxiliary Variables. There is a problem with the specification of Fibonacci that we have been using as a running example. The specification ($n > 0, x == \text{Fib}(n)$) has trivial solutions, since nothing prevents the program from modifying the value of the input variable n . Thus the triple $\{n > 0\} n := 0 ; x := 0 \{x == \text{Fib}(n)\}$ is a valid Hoare triple.

The problem can be avoided with the use of an *auxiliary variable* to record the value of n in the pre-state. The triple

$$\{n > 0 \ \&\& \ n == n_0\} F \{x == \text{Fib}(n) \ \&\& \ n == n_0\}$$

solves the problem as long as n_0 is indeed auxiliary, i.e. it is not used as a program variable in F . Hoare logic has no explicit support for auxiliary variables: it is not possible to write the specification above in a way that formally prevents n_0 from being modified by F . This is particularly relevant if one is interested in doing modular verification, since one may have to reason about a program that calls an external procedure for which only the specification, and not the code, is available (see Section 10.3 below).

In the context of program verification tools, a common alternative to auxiliary variables consists in enriching the assertion language with an operator allowing one to refer to the value of an expression at a given program point.

For instance, our example specification could be written

$$\{n > 0\} F \{x == \text{Fib}(n) \ \&\& \ n == \text{old}(n)\}$$

or simply

$$\{n > 0\} F \{x == \text{Fib}(\text{old}(n))\}$$

where $\text{old}(n)$ refers to the value of n in the initial state in which F is executed.

The generation of verification conditions can be altered to cope with this operator by first substituting, for every occurrence of the operator in the postcondition, x_0 for $\text{old}(x)$, where x_0 is a fresh auxiliary variable, and at the same time strengthening the precondition with the equality formula $x == x_0$.

8 Guarded Commands

We take a detour here to briefly consider a variant of Dijkstra's guarded commands language and *Weakest Precondition Calculus* (WPC). The calculus provides a verification conditions generator, and guarded commands, although apparently very abstract, are used as intermediate language by at least two standard program verification tools.

The language is different from the programming language we have been considering in a number of ways. First, there is no distinction between Boolean expressions and assertions. The language contains two primitives that test the value of assertions: the commands **assert** b and **assume** b both behave like **skip** if b evaluates to *true*. The difference is that the former *terminates abruptly* if b evaluates to *false*, whereas the latter *cannot be executed*. It is thus a partial command that can be used as a *guard* for the execution of a subsequent command.

Another aspect is the presence of a *non-deterministic choice* operator. The command $C_1 \parallel C_2$ will arbitrarily execute either C_1 or C_2 .

The expression $wp.C.Q$ denotes the weakest precondition such that Q holds as a postcondition if the command C terminates. The following calculus defines the weakest precondition semantics of the language.

$$\begin{aligned} wp.(\mathbf{assert} \ b).Q &= b \ \&\& \ Q \\ wp.(\mathbf{assume} \ b).Q &= b \rightarrow Q \\ wp.(x := e).Q &= Q[x \mapsto e] \\ wp.(C_1; C_2).Q &= wp.C_1.(wp.C_2.Q) \\ wp.(C_1 \parallel C_2).Q &= wp.C_1.Q \ \&\& \ wp.C_2.Q \end{aligned}$$

Given a program C , the verification condition generated for the partial correctness Hoare triple $\{P\} C \{Q\}$ is given by

$$\text{VC}(\{P\} C \{Q\}) = [P \rightarrow wp.C.Q]$$

In fact, since it corresponds to partial correctness, this notion is usually known as the weakest *liberal* precondition.

Two crucial constructs of an imperative language – conditionals and loops – are missing from the guarded commands language, but they can be encoded, i.e. translated into guarded commands that generate the appropriate weakest preconditions. The command **if** b **then** C_t **else** C_f can be encoded as

$$(\text{assume } b; C_t) \parallel (\text{assume } !b; C_f)$$

where b and $!b$ are used as guards in a choice command, which has the effect of removing non-determinism since it is certain that one of the commands cannot be executed.

Loops can be encoded in a number of ways. If one is willing to give up on soundness and check that the invariant holds for only a limited number of iterations, the loop **while** b **do** $\{I\} C$ can simply be encoded as follows

$$\begin{aligned} & \text{assert } I; \\ & (\text{assume } b; C; \text{assert } I; \text{assume false}) \parallel (\text{assume } !b) \end{aligned}$$

Here the loop is taken to execute at most once but this can easily be expanded to an arbitrary fixed number of iterations. The effect of the first **assert** command is to test the invariant in the initial conditions; the first branch of the choice command tests its preservation by an iteration of the loop body; the second branch establishes the falsity of the loop condition on exit. Note that if the choice selects the first branch, the second branch will inevitably be selected when **assume false** is reached (the command in the first branch cannot be executed in its entirety), thus this encodes a loop that iterates at most once.

To test the preservation of the invariant in an *arbitrary* iteration of the loop requires to identify all the variables (say x_1, \dots, x_n) assigned in the loop body, and to assign them arbitrary values. Fresh variables (y_1, \dots, y_n) can be used to this effect. The following translation of the loop first tests that the invariant is initially true, then resets the values of variables, and assumes the truth of the invariant for the current arbitrary state. The choice command that follows is the same as before, but the first branch is now testing the preservation of

the invariant in an arbitrary iteration.

```
assert  $I$  ;  
 $x_1 := y_1 ; \dots ; x_n := y_n$  ;  
assume  $I$  ;  
(assume  $b$  ;  $C$  ; assert  $I$  ; assume false)  $\square$  (assume  $!b$ )
```

The specific guarded command language presented in this section has been widely used as a core language in the development of at least two major tools for checking the behaviour of programs: the ESC [65, 63, 17] family of tools (of which ESC/Modula3, ESC/Java and ESC/Java2 are instances) and more recently Boogie [5], a generic VCGen that is being used notably with the Spec# language. Both tools are capable of generating verification conditions as proof obligations for the Simplify [28] prover, but Boogie supports more recent and advanced proof tools.

ESC stands for “extended static checking”; its emphasis was more on providing programmers with tools that could find common errors (such as null dereferencing) rather than on program verification. Boogie is a very sophisticated tool that integrates advanced features such as automatic inference of loop invariants. It is a generic VCGen in the sense that different programming languages can be translated into the BoogiePL language (in a sound way with respect to the generated verification conditions). VC generation based on weakest preconditions has been advanced with the development of Boogie to cover, for instance, programs containing both loops and **goto** statements [6].

Both tools can be used with complex annotation languages for real-world programs. For instance ESC/Java uses JML (a standard annotation language for Java programs, see Section 11), and Boogie has been used as a VCGen for Spec# (similar to JML but for C# programs). In both tools the guarded command language is used as an intermediate language into which source code is translated; verification conditions are generated by applying the weakest-precondition calculus.

An important result was proposed as part of the development of ESC. The simple definition of weakest precondition given above generates VCs whose size is potentially exponential in the size of the source code. Two clauses in the definition of wp are responsible for this: the case of the assignment command $x := e$, which may have to create as many copies of the expression e as there are occurrences of x in the postcondition Q ; and the choice command, which duplicates Q .

This problem can be fixed [37, 64] by using a two-stage algorithm that pro-

duces VCs that are worst-case quadratic in size (and usually close to linear). The idea is that guarded commands are first translated into a *passive form*, where assignments are eliminated, replaced by *assume* commands and fresh variables for each assigned variable. For instance, $x := x + 1; \mathbf{assert} x > 0$ becomes $\mathbf{assume} x_1 = x + 1; \mathbf{assert} x_1 > 0$. This translation preserves the weakest precondition semantics, and while it may increase the size of the code, this growth is worst-case quadratic and near linear in practice.

Passive commands may not affect the state of programs since they do not contain assignment statements. The semantics of two programs may differ only with respect to termination: programs may terminate normally or erroneously. The weakest preconditions $wp.C.true$ and $wp.C.false$ characterize respectively the states from which C may not terminate erroneously and the states from which C may not terminate normally. The weakest precondition for a given postcondition Q can then be written as

$$wp.C.Q = wp.C.true \ \&\& \ (!wp.C.false \rightarrow Q)$$

The resulting condition is worst-case quadratic in the size of the code.

The weakest precondition technique is quite flexible. For instance it is easy to treat erroneous termination explicitly by calculating weakest preconditions with respect to two postconditions. The following extended definition ensures a given postcondition Q if the program terminates normally and a possibly distinct postcondition R if it terminates abruptly with a failed **assert** command.

$$\begin{aligned} wp.(\mathbf{assert} b).(Q, R) &= (b \ \&\& \ Q) \ || \ (!b \ \&\& \ R) \\ wp.(\mathbf{assume} b).(Q, R) &= b \rightarrow Q \\ wp.(x := e).(Q, R) &= Q[x \mapsto e] \\ wp.(C_1; C_2).(Q, R) &= wp.C_1.(wp.C_2.(Q, R), R) \\ wp.(C_1 \ || \ C_2).(Q, R) &= wp.C_1.(Q, R) \ \&\& \ wp.C_2.(Q, R) \end{aligned}$$

Exceptional termination can also be considered in this framework, which inspires our treatment of exceptions for the While language in Section 10.1.

The weakest preconditions calculus has also been studied extensively and applied in many textbooks [29, 55, 42] as a tool for the development of *correct-by-construction* software – quite a different perspective from program verification.

9 Hoare Logic with Updates

In Section 6 it was shown that the intermediate assertions required by the sequencing rule of Hoare logic could be calculated either backwards from post-conditions or in a forward fashion starting from the preconditions. In order to define an inference system that propagates assertions in a forward manner, it is convenient to modify the abstract syntax of programs to the following linear version, in which programs are defined as (possibly empty) sequences of commands. We let **PROG** denote the class of programs.

$$\mathbf{PROG} \ni W ::= C ; W \mid \varepsilon$$

$$\mathbf{ICOMM} \ni C ::= \mathbf{skip} \mid x := e \mid \mathbf{if } b \mathbf{ then } W \mathbf{ else } W \mid \mathbf{while } b \mathbf{ do } \{A\} W$$

This corresponds simply to a left-associative view of the sequence operator. We define a translation function $\mathcal{T} : \mathbf{IComm} \rightarrow \mathbf{PROG}$ as follows

$$\begin{aligned} \mathcal{T}(\mathbf{skip}) &= \mathbf{skip} ; \varepsilon \\ \mathcal{T}(x := e) &= x := e ; \varepsilon \\ \mathcal{T}(C_1 ; C_2) &= \mathcal{T}(C_1) @ \mathcal{T}(C_2) \\ \mathcal{T}(\mathbf{if } b \mathbf{ then } C_t \mathbf{ else } C_f) &= \mathbf{if } b \mathbf{ then } \mathcal{T}(C_t) \mathbf{ else } \mathcal{T}(C_f) ; \varepsilon \\ \mathcal{T}(\mathbf{while } b \mathbf{ do } \{I\} C) &= \mathbf{while } b \mathbf{ do } \{I\} \mathcal{T}(C) ; \varepsilon \end{aligned}$$

where $\cdot @ \cdot$ denotes sequence concatenation.

A natural semantics for programs and commands can be given via a mutually dependent definition of the evaluation relation, similar to what is done in Figure 2 (with the obvious adaptations).

The formulation of specifications that will be defined for these programs uses a notion of *update*. An update is simply a partial mapping from variables to expressions of the language. We write $\{x_1 \mapsto e_1, \dots, x_n \mapsto e_n\}$ for the update mapping x_i to e_i , for $i \in \{1, \dots, n\}$, and \emptyset for the empty update. Let $\mathcal{U} = \{x_1 \mapsto e_1, \dots, x_n \mapsto e_n\}$ be an update, and t an expression (resp. assertion). We write $\mathcal{U}(t)$ to denote the expression (resp. assertion) obtained by the simultaneous substitution of terms e_i for the free occurrences of variables x_i in t .

Updates are modified by an operation that sets the value of a variable, defined as

$$(\mathcal{U} ; x := e) = \mathcal{U} \oplus \{x \mapsto \mathcal{U}(e)\}$$

where \oplus denotes function overriding. For two partial functions $f, g : X \leftrightarrow Y$ we define $f \oplus g : X \leftrightarrow Y$ as follows

$$(f \oplus g)(x) = \begin{cases} g(x) & \text{if } x \in \text{dom}(g) \\ f(x) & \text{if } x \notin \text{dom}(g) \wedge x \in \text{dom}(f) \end{cases}$$

The notion of specification will now be slightly modified to correspond to a Hoare triple extended with an update, written $\{P\}[\mathcal{U}]W\{Q\}$. The idea is that the update is applied to the initial state, in which the precondition holds, before the program is executed. The semantic interpretation of such a specification is

$$\llbracket \{P\}[\mathcal{U}]W\{Q\} \rrbracket = \forall s, s' \in \Sigma. \llbracket P \rrbracket(s) \wedge (W, s_{\mathcal{U}}) \Downarrow s' \Rightarrow \llbracket Q \rrbracket(s')$$

where the updated state $s_{\mathcal{U}}$ is defined as follows:

$$s_{\mathcal{U}}(x) = \begin{cases} s(x) & \text{if } x \notin \text{dom}(\mathcal{U}) \\ \llbracket \mathcal{U}(x) \rrbracket(s) & \text{if } x \in \text{dom}(\mathcal{U}) \end{cases}$$

Before introducing the inference system we proceed with some lemmas involving the concepts described.

Lemma 17 *For all $C \in \mathbf{IComm}$, $(\mathcal{T}(C), s) \Downarrow s'$ iff $(C, s) \Downarrow s'$*

Proof. By induction on the structure of C . □

Proposition 18 *For $C \in \mathbf{IComm}$, $\llbracket \{P\}[\emptyset]\mathcal{T}(C)\{Q\} \rrbracket = \llbracket \{P\}C\{Q\} \rrbracket$*

Proof. Immediate by Lemma 17. □

Lemma 19

$(W_1 @ W_2, s) \Downarrow s'$ iff $(W_1, s) \Downarrow s''$ and $(W_2, s'') \Downarrow s'$, for some $s'' \in \Sigma$

Proof. By induction on the length of W_1 . □

Lemma 20 $\llbracket \mathcal{U}(Q) \rrbracket(s) = \llbracket Q \rrbracket(s_{\mathcal{U}})$

Proof. By induction on the structure of Q . □

Lemma 21 (1) $s_{(\mathcal{U}; x:=e)} = (s_{\mathcal{U}})_{x:=e}$

(2) $(W, s_{(\mathcal{U}; x:=e)}) \Downarrow s'$ iff $(x := e; W, s_{\mathcal{U}}) \Downarrow s'$

Proof. Directly by unfolding the definitions. □

The inference system of Hoare logic with updates is given in Figure 15. We name it system **Hu**. There is one rule for each program construct, except for

$$\begin{array}{c}
\frac{}{\{P\}[\mathcal{U}]\varepsilon\{Q\}} \quad \text{if } \models P \rightarrow \mathcal{U}(Q) \\
\\
\frac{\{P\}[\mathcal{U}]W\{Q\}}{\{P\}[\mathcal{U}]\text{skip}; W\{Q\}} \qquad \frac{\{P\}[\mathcal{U}; x := e]W\{Q\}}{\{P\}[\mathcal{U}]x := e; W\{Q\}} \\
\\
\frac{\{I \ \&\& \ b\}[\emptyset]W_t\{I\} \quad \{I \ \&\& \ !b\}[\emptyset]W\{Q\}}{\{P\}[\mathcal{U}]\text{while } b \text{ do } \{I\}W_t; W\{Q\}} \quad \text{if } \models P \rightarrow \mathcal{U}(I) \\
\\
\frac{\{P \ \&\& \ \mathcal{U}(b)\}[\mathcal{U}]W_t@W\{Q\} \quad \{P \ \&\& \ !\mathcal{U}(b)\}[\mathcal{U}]W_f@W\{Q\}}{\{P\}[\mathcal{U}]\text{if } b \text{ then } W_t \text{ else } W_f; W\{Q\}}
\end{array}$$

Fig. 15. Rules of Hoare logic with updates – System **Hu**

sequencing, since all programs are now seen as sequences. The rule for the empty program (called the *exit* rule) introduces a verification condition; the remaining rules are selected by pattern-matching on the first command of the program. So, the construction of derivation trees is syntax-directed.

Proposition 22 (Soundness) *Every specification derivable in system **Hu** is semantically valid: If $\vdash_{\mathbf{Hu}} \{P\}[\mathcal{U}]W\{Q\}$, then $\llbracket \{P\}[\mathcal{U}]W\{Q\} \rrbracket = \text{true}$.*

Proof. By induction on the derivation of $\vdash_{\mathbf{Hu}} \{P\}[\mathcal{U}]W\{Q\}$. □

So, it follows directly from propositions 18 and 22 that to prove the validity of the triple $\{P\}C\{Q\}$ it is enough to show that $\vdash_{\mathbf{Hu}} \{P\}[\emptyset]\mathcal{T}(C)\{Q\}$.

Applying the rules of system **Hu** backwards to construct a proof tree is in fact very close to performing a symbolic execution of the program, since at each step of the proof, corresponding to a program $C; W$, the proof construction will proceed with the execution of the program W .

For this reason it has been argued that tools based on such a system with updates may be more adequate for debugging purposes, since the construction of proofs follows the symbolic execution of the code.

Figure 16 contains the definition of a VCGen obtained from system **Hu**. We remark that whereas in the case of system **Hgi** producing a VCGen required imposing a specific strategy for the construction of proof trees, in system **Hu** the construction of derivations is deterministic. It is immediate to see that the verification conditions produced by this algorithm, for a given Hoare triple with updates, are exactly the side conditions of its derivation tree.

$$\begin{aligned}
\text{VCGu}(\{P\}[\mathcal{U}] \varepsilon \{Q\}) &= \{[P \rightarrow \mathcal{U}(Q)]\} \\
\text{VCGu}(\{P\}[\mathcal{U}] \text{skip}; W \{Q\}) &= \text{VCGu}(\{P\}[\mathcal{U}] W \{Q\}) \\
\text{VCGu}(\{P\}[\mathcal{U}] x := e; W \{Q\}) &= \text{VCGu}(\{P\}[\mathcal{U}; x := e] W \{Q\}) \\
\text{VCGu}(\{P\}[\mathcal{U}] \text{while } b \text{ do } \{I\} W_t; W \{Q\}) &= \{[P \rightarrow \mathcal{U}(I)]\} \cup \text{VCGu}(\{I \ \&\& \ b\}[\emptyset] W_t \{I\}) \cup \\
&\quad \text{VCGu}(\{I \ \&\& \ !b\}[\emptyset] W \{Q\}) \\
\text{VCGu}(\{P\}[\mathcal{U}] \text{if } b \text{ then } W_t \text{ else } W_f; W \{Q\}) &= \text{VCGu}(\{P \ \&\& \ \mathcal{U}(b)\}[\mathcal{U}] W_t @ W \{Q\}) \cup \\
&\quad \text{VCGu}(\{P \ \&\& \ !\mathcal{U}(b)\}[\mathcal{U}] W_f @ W \{Q\})
\end{aligned}$$

Fig. 16. A VCGen based on updates: VCGu

$$\begin{aligned}
&\text{VCGu}(\{P\}[\emptyset] C_{\text{Fib}}^I; \varepsilon \{Q\}) \\
&= \text{VCGu}(\{P\}[\{x \mapsto 1, y \mapsto 0, i \mapsto 1\}] \text{while } i < n \text{ do } \{I\} C''; \varepsilon \{Q\}) \\
&= \{[P \rightarrow \{x \mapsto 1, y \mapsto 0, i \mapsto 1\}(I)]\} \cup \text{VCGu}(\{I \ \&\& \ i < n\}[\emptyset] C'' \{I\}) \cup \text{VCGu}(\{I \ \&\& \ !(i < n)\}[\emptyset] \varepsilon \{Q\}) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)]\} \cup \{[I \ \&\& \ !(i < n) \rightarrow \emptyset(Q)]\} \cup \\
&\quad \text{VCGu}(\{I \ \&\& \ i < n\}[\emptyset] aux := y; y := x; x := x + aux; i := i + 1; \varepsilon \{I\}) \\
&= \{[n > 0 \rightarrow 1 \leq n \ \&\& \ 1 == \text{Fib}(1) \ \&\& \ 0 == \text{Fib}(1-1)], \\
&\quad [i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1) \ \&\& \ !(i < n) \rightarrow Q], \\
&\quad [i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1) \ \&\& \ i < n \rightarrow \\
&\quad \quad i + 1 \leq n \ \&\& \ x + y == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)]\} \\
&\text{since,} \\
&\text{VCGu}(\{I \ \&\& \ i < n\}[\emptyset] aux := y; y := x; x := x + aux; i := i + 1; \varepsilon \{I\}) \\
&= \text{VCGu}(\{I \ \&\& \ i < n\}[\emptyset; aux := y] y := x; x := x + aux; i := i + 1; \varepsilon \{I\}) \\
&= \text{VCGu}(\{I \ \&\& \ i < n\}[\{aux \mapsto y\}; y := x] x := x + aux; i := i + 1; \varepsilon \{I\}) \\
&= \text{VCGu}(\{I \ \&\& \ i < n\}[\{aux \mapsto y, y \mapsto x\}; x := x + aux] i := i + 1; \varepsilon \{I\}) \\
&= \text{VCGu}(\{I \ \&\& \ i < n\}[\{aux \mapsto y, y \mapsto x, x \mapsto x + y\}; i := i + 1] \varepsilon \{I\}) \\
&= \{[I \ \&\& \ i < n \rightarrow \{aux \mapsto y, y \mapsto x, x \mapsto x + y, i \mapsto i + 1\}(I)]\} \\
&= \{[i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1) \ \&\& \ i < n \rightarrow \\
&\quad i + 1 \leq n \ \&\& \ x + y == \text{Fib}(i+1) \ \&\& \ x == \text{Fib}((i+1)-1)]\}
\end{aligned}$$

Fig. 17. Fibonacci example, using VCGu

Lemma 23 $\text{VCGu}(\{P\}[\mathcal{U}] W \{Q\}) \Vdash_{\text{Hu}} \{P\}[\mathcal{U}] W \{Q\}$

Proof. By induction on the structure of W . □

Example 24 We exemplify the use of this VCGen with our running example of Figure 11. Let P be $n > 0$, Q be $x == \text{Fib}(n)$, and I be $i \leq n \ \&\& \ x == \text{Fib}(i) \ \&\& \ y == \text{Fib}(i-1)$. We start with the empty update. The set of generated verification conditions is shown in Figure 17.

Hoare logic with updates was introduced in [43] with the accompanying proof tool KeY-Hoare. The tool reads in a specification and try to prove it by constructing a proof tree (no explicit VCGen is used). The tool also includes a first-order theory.

KeY-Hoare is in fact a side-product of a much bigger effort. The full-fledged KeY tool [1] is a verification tool for Java Card programs (compatible with JML annotations) capable of handling many aspects of real-world object-oriented programs. KeY is based on JavaCardDL [10], a version of dynamic logic [44] suited for this language. KeY also includes a symbolic debugger module.

Updates are a key ingredient of JavaCardDL; the other main ingredient, which stands at the heart of dynamic logic, is the existence of *modalities* in the assertion language; in particular, for each program W in the programming language and assertion Q , there exists a new assertion $[W]Q$, interpreted informally as “if execution of W terminates, then Q will hold in the final state”.

Formulas of JavaCardDL are of the form $([W]Q)^u$. This formula is interpreted as *true* in a state s if Q holds in the state that results from executing W in the state s_u , if execution terminates.

The Hoare triple $\{P\}C\{Q\}$ can be written as the formula $P \rightarrow [C]Q$, and it is not difficult to see that system **Hu** can be rewritten in the syntax of dynamic logic. Note however that, since assertions of dynamic logic may contain programs, it may not be possible to interpret an arbitrary formula like $P \rightarrow [C]Q$ directly as a Hoare triple.

10 Language Extensions

In this section we discuss several extensions to the initial language. Our aim here is to provide some insight into how each of these aspects is treated by common program verification tools for real-world languages. A brief review of these tools is then given in the final section of the paper.

10.1 Exceptions

Adding exceptions to our language is useful not only because an exception mechanism is useful in itself, but also because it provides a means to model control-transfer commands (like **break** in C and Java).

We extend the syntax with two new commands, and a new form of specifications for exceptional termination.

$$\begin{aligned} \mathbf{IComm} \ni C &::= \dots \mid \mathbf{try} C \mathbf{catch} C \mid \mathbf{throw} \\ \mathbf{Spec} \in S &::= \{A\} C \{A\} \mid \{A\} C \{\!\!|A\!\!\} \end{aligned}$$

The **throw** command raises an exception (execution terminates abruptly). The command **try** C **catch** C_c executes C and catches a possible exception raised by it, in which case the code C_c is executed.

The informal meaning of a specification $\{P\} C \{\!\!|Q\!\!\}$ is that if the program C is executed in an initial state in which the precondition P is *true*, then either execution of C does not terminate or it terminates *with an exception raised*, in which case the postcondition Q is *true* in the final state. If the triple $\{P\} C \{Q\}$ is valid and C terminates, then it does so normally, with no exception raised.

The behaviour of **throw** and **try** C **catch** C_c can be described axiomatically by rules like the following, where in the latter case one rule is given for normal termination of C and another for exceptional termination.

$$\begin{array}{c} \frac{}{\{P\} \mathbf{throw} \{\!\!|Q\!\!\}} \quad \text{if } \models P \rightarrow Q \qquad \frac{\{P\} C \{Q\}}{\{P\} \mathbf{try} C \mathbf{catch} C_c \{Q\}} \\ \\ \frac{\{P\} C \{\!\!|S\!\!\} \quad \{S\} C_c \{Q\}}{\{P\} \mathbf{try} C \mathbf{catch} C_c \{Q\}} \qquad \frac{\{P\} C \{\!\!|S\!\!\} \quad \{S\} C_c \{\!\!|R\!\!\}}{\{P\} \mathbf{try} C \mathbf{catch} C_c \{\!\!|R\!\!\}} \end{array}$$

We remark that simply adding these rules to the system **Hgi** would not work: one would also have to add new rules, similar to those of **Hgi**, to account for exceptional termination. This presentation is a bit tedious.

A more compact alternative is to use a special notation for specifications with two possible outcomes. The specification $\{P\} C \{Q\}\{\!\!|R\!\!\}$ is valid if either $\{P\} C \{Q\}$ is valid or $\{P\} C \{\!\!|R\!\!\}$ is valid. With this notion of specification it is possible to write a goal-directed Hoare logic system for the extended language by simply adding an exceptional postcondition $\{\!\!|R\!\!\}$ to all the Hoare triples in the rules of system **Hgi**, and extending it with the following rules

$$\frac{}{\{P\} \mathbf{throw} \{Q\}\{\!\!|R\!\!\}} \quad \text{if } \models P \rightarrow R \qquad \frac{\{P\} C \{Q\}\{\!\!|S\!\!\} \quad \{S\} C_c \{Q\}\{\!\!|R\!\!\}}{\{P\} \mathbf{try} C \mathbf{catch} C_c \{Q\}\{\!\!|R\!\!\}}$$

The definitions of **wprec** and **VCG** (figures 9 and 13) can be easily adapted for this extended language; it suffices to modify the signatures of functions to take two postconditions as arguments, including a second postcondition R

in every invocation $\text{wprec}(\cdot, \cdot, R)$ and $\text{VC}_{\text{aux}}(\cdot, \cdot, R)$. The following additional clauses are required.

$$\text{wprec}(\text{throw}, Q, R) = R$$

$$\text{wprec}(\text{try } C \text{ catch } C_c, Q, R) = \text{wprec}(C, Q, \text{wprec}(C_c, Q, R))$$

$$\text{VC}_{\text{aux}}(\text{throw}, Q, R) = \emptyset$$

$$\text{VC}_{\text{aux}}(\text{try } C \text{ catch } C_c, Q, R) = \text{VC}_{\text{aux}}(C, Q, \text{wprec}(C_c, Q, R)) \cup \text{VC}_{\text{aux}}(C_c, Q, R)$$

And finally

$$\text{VCG}(\{P\} C \{Q\} \{R\}) = \{[P \rightarrow \text{wprec}(C, Q, R)]\} \cup \text{VC}_{\text{aux}}(C, Q, R)$$

This treatment of exceptions is similar to that proposed in the work of Leino and colleagues on weakest preconditions of guarded command languages [65].

10.2 Arrays and Pointers: Aliasing

Aliasing is a phenomenon that occurs in programming languages when some form of indirection is possible.

Arrays provide one such example: if we introduce arrays of integers in our language (with notation $u[k]$ for the value stored in position k of array u), then of course indexing by arbitrary expressions must be allowed – restricting indexes to constants makes arrays useless collections of variables, since it is not possible to iterate over them. This creates an opportunity for aliasing to occur. In particular, the expressions $u[e]$ and $u[e']$ may refer to the same positions in the array or not, depending on whether $e = e'$ holds. This is usually called *subscript aliasing*.

Simply viewing arrays as indexed variables will not work. Consider the following adaptation of the weakest precondition calculation for an assignment instruction.

$$\text{wprec}(u[i] := e, Q) = Q[u[i] \mapsto e]$$

This would yield for instance

$$\text{wprec}(u[i] := 10, u[j] > 100) = u[j] > 100$$

but clearly the precondition $u[j] > 100$ will not be preserved by the command, if executed in a state in which $i = j$.

The correct output would be a logical condition that incorporates all the required comparisons between all index expressions present in the postcondition and the array position assigned to. For the above example one would have

$$\text{wprec}(u[i] := 10, u[j] > 100) = (i == j \rightarrow 10 > 100) \ \&\& \ (i != j \rightarrow u[j] > 100)$$

It is relatively simple to devise an algorithm to produce this verification condition.

Hoare's solution to this problem was slightly different: arrays were seen as monolithic objects, equipped with an array update operation. Let $\text{write}(u, i, e)$ denote the array that results from updating u by setting the contents of position i to e ; then the weakest precondition can be calculated correctly as

$$\text{wprec}(u[i] := e, Q) = Q[u \mapsto \text{write}(u, i, e)]$$

This approach involves reasoning in the context of what is usually described as a *theory of arrays*, which specifies the behaviour of the update operation. The two most fundamental axioms are

$$\text{Forall } u, i, j, e. i == j \rightarrow \text{write}(u, i, e)[j] == e \quad (1)$$

$$\text{Forall } u, i, j, e. i \neq j \rightarrow \text{write}(u, i, e)[j] == u[j] \quad (2)$$

Proof tools aimed at program verification, in particular SMT solvers, typically allow reasoning with such an array theory.

The logic with updates studied in Section 9 copes very easily with aliasing: it suffices to adapt to array positions or structure fields the definition of the operation that sets the value of a variable. For instance, for array positions one could write

$$(\mathcal{U}; a[j] := e) = \mathcal{U} \oplus \{a[j] \mapsto \mathcal{U}(e)\}$$

A different form of aliasing occurs in a language with structures (or objects). Let $s.a$ denote the value stored in the field a of structure s ; then $p.a$ and $q.a$ may refer to the same value or not, depending on whether p and q refer to the same structure. Note that this only makes sense if p and q are *references* (or pointers) to structures, since it is not possible for two ordinary variables to have as values the same structure. For this reason this form of aliasing is known as *pointer aliasing*.

Dealing with pointer aliasing is crucial if one wants to be able to verify programs that use recursive data structures, dynamically defined in heap memory. A heap can be seen as a very big array (indexed by memory positions), and in this sense pointer aliasing is an instance of index aliasing. Seriously reasoning about pointer programs requires a new framework that avoids the proliferation of arithmetic proof obligations concerning indexes. The theoretical advances

in this area have been developed in the context of *separation logic* [73, 77], whose key idea is the possibility to explicitly express the separation between different structures.

An alternative to this approach had previously been proposed by Burstall [16] and further explored by Bornat [14]. The idea here was to see the heap as a *collection* of arrays, rather than a single array. In particular, there should be one such array for every structure/object field. For instance, $p.a$ and $q.a$ could be represented by positions $a[p]$ and $a[q]$ in the same (heap model) array a . They will be represented by the same position if p and q are the same memory address. The operation of setting the value of field a of the structure referred by p to e is modelled by `write(a, p, e)`.

This approach has the advantage of producing first-order verification conditions, which means that (unlike separation logic) it can be handled by a standard prover in a straightforward fashion, requiring simply a theory of arrays. The Caduceus VCGen uses this approach to construct a memory model for C programs. The heap model is explicitly constructed as a set of data structures in the language of the Why generic VCGen (see Section 11 below).

10.3 Procedure Calls

The inclusion of procedures in the programming language raises a number of issues that will now be briefly discussed. For the sake of simplicity we restrict our attention to parameterless procedures.

Let the code $C_{\mathbf{p}}$ be the body of procedure \mathbf{p} ; the command `call \mathbf{p}` invokes this procedure and transfers control to $C_{\mathbf{p}}$. After execution of $C_{\mathbf{p}}$, control will return to the point immediately after `call \mathbf{p}` in the caller. The following straightforward rule states that reasoning about `call \mathbf{p}` should be transferred to reasoning about $C_{\mathbf{p}}$.

$$\frac{\{P\} C_{\mathbf{p}} \{Q\}}{\{P\} \text{call } \mathbf{p} \{Q\}}$$

The first difficulty is that if $C_{\mathbf{p}}$ contains a call to \mathbf{p} this will generate infinite derivations. From the point of view of verification conditions we could define the weakest precondition $\text{wprec}(\text{call } \mathbf{p}, Q) = \text{wprec}(C_{\mathbf{p}}, Q)$, but expanding this would of course lead to a recursive equation $\text{wprec}(\text{call } \mathbf{p}, Q) = P'$ where P' contains occurrences of $\text{wprec}(\text{call } \mathbf{p}, Q')$ for some assertions Q' .

Hoare [46] proposed the following rule to reason about procedure calls in the

presence of recursion:

$$\frac{\begin{array}{c} [\{P\} \text{call } \mathbf{p} \{Q\}] \\ \vdots \\ \{P\} C_{\mathbf{p}} \{Q\} \end{array}}{\{P\} \text{call } \mathbf{p} \{Q\}}$$

The rule makes use of the notion of hypothetical derivation, as in the inference systems for natural deduction. If reasoning under the assumption $\{P\} \text{call } \mathbf{p} \{Q\}$, one is able to prove that $\{P\} C_{\mathbf{p}} \{Q\}$ holds, this amounts in fact to proving the assumption itself, which can be *cancelled* in the derivation. Intuitively, if assuming the correctness of \mathbf{p} wrt. (P, Q) , one is able to prove that same correctness, then \mathbf{p} is indeed correct wrt. (P, Q) .

A second difficulty now arises: this rule is not appropriate to extend system Hgi . Observe that while reasoning about $C_{\mathbf{p}}$ one will have to use as hypothesis the triple $\{P\} \text{call } \mathbf{p} \{Q\}$, and system Hgi does not provide a consequence rule to adapt P and Q as required by the structure of $C_{\mathbf{p}}$. Additionally, it is also not clear how verification conditions can be generated following this rule.

Similarly to reasoning about loops, this difficulty can be addressed by introducing additional annotations in the code. These annotations will consist of a precondition and a postcondition for each procedure, usually known as the procedure's *contract*. In the following, whenever we say that a procedure is correct, it is meant that it is correct with respect to its contract.

Let then precondition $P_{\mathbf{p}}$ and postcondition $Q_{\mathbf{p}}$ be the contract or public specification of \mathbf{p} , and consider for a moment that no auxiliary variables are allowed to occur in $P_{\mathbf{p}}$ and $Q_{\mathbf{p}}$. The following two rules can be added to system Hgi for reasoning about procedures.

$$\begin{array}{l} \text{(contract)} \\ \text{(adapt)} \end{array} \quad \frac{\begin{array}{c} [\{P_{\mathbf{p}}\} \text{call } \mathbf{p} \{Q_{\mathbf{p}}\}] \\ \vdots \\ \{P_{\mathbf{p}}\} C_{\mathbf{p}} \{Q_{\mathbf{p}}\} \end{array}}{\{P_{\mathbf{p}}\} \text{call } \mathbf{p} \{Q_{\mathbf{p}}\}} \quad \frac{\{P_{\mathbf{p}}\} \text{call } \mathbf{p} \{Q_{\mathbf{p}}\}}{\{P\} \text{call } \mathbf{p} \{Q\}} \quad \text{if } \models P \rightarrow P_{\mathbf{p}} \text{ and } \models Q_{\mathbf{p}} \rightarrow Q$$

The first rule is used to establish the correctness of the procedure with respect to its contract only, not with respect to other specifications. The second rule is a dedicated version of the consequence rule, usually known as an *adaptation* rule, which can only be used to prove that \mathbf{p} meets any specification that is weaker than its contract. Crucially, unlike a general consequence rule, this rule enjoys the subformula property, since $P_{\mathbf{p}}$ and $Q_{\mathbf{p}}$ are obtained from \mathbf{p} . The adaptation rule can be used notably in the derivation of the triple

$\{P_{\mathbf{p}}\} \mathbf{call} \mathbf{p} \{Q_{\mathbf{p}}\}$ using the (*contract*) rule, to adapt the assumption triple to the precondition and postcondition that may be locally required.

Before considering how verification conditions can be generated based on this rule, consider now that we have a set of procedures that invoke each other freely in a mutually recursive way; a *program* can be seen as a set of such mutually recursive procedures (a similar context is that of a *class* in an object oriented programming language, consisting of a set of methods that share a number of class attributes, playing the role of global variables). To formalise this setting, we assume a set of procedure names \mathcal{P} and let \mathbf{p} range over \mathcal{P} . Let **Proc** denote the class of procedure definitions (the body of a procedure definition is just a specification). We extend the command syntax with procedure calls, and let programs be defined as sequences of procedure definitions.

$$\begin{aligned} \mathbf{IComm} \ni C &::= \dots \mid \mathbf{call} \mathbf{p} \\ \mathbf{Proc} \ni \Phi &::= \mathbf{proc} \mathbf{p} = \{A\} C \{A\} \\ \mathbf{Prog} \ni \Pi &::= \Phi \Pi \mid \Phi \end{aligned}$$

The procedure definition $\mathbf{proc} \mathbf{p} = \{P_{\mathbf{p}}\} C_{\mathbf{p}} \{Q_{\mathbf{p}}\}$ sets the public specification of \mathbf{p} to consist of precondition $P_{\mathbf{p}}$ and postcondition $Q_{\mathbf{p}}$, and the body of \mathbf{p} to be $C_{\mathbf{p}}$. We say that a program Π is well-defined if the names of its procedures are unique, and no procedure in Π invokes procedures which are not in Π . One particular procedure should be designated as an entry point into each program, but we will abstract away from this operational issue.

The word ‘contract’ is used in the sense that the specification of each procedure establishes an agreement between a calling procedure and the invoked procedure: if \mathbf{p}_1 calls \mathbf{p}_2 then \mathbf{p}_1 should ensure that the precondition of \mathbf{p}_2 is satisfied immediately before the call; the specification of \mathbf{p}_2 guarantees that its postcondition will be satisfied when control is returned to \mathbf{p}_1 . This principle is thoroughly explored in the so-called *design-by-contract* approach to software development [69].

How should the (*contract*) rule be modified to cope with multiple procedures and mutual recursion? The rationale must now be “if assuming the correctness of all procedures in a program, one is able to prove individually the correctness of every procedure in it, then all procedures are correct”. It is not possible in general to assert individually the correctness of a single procedure with an empty set of assumptions. Let Π be a program and \mathbf{p}_i range over its procedures,

with $i \in \{1, \dots, n\}$. The rules now become:

$$\begin{array}{c}
\text{(contract)} \\
\frac{
\begin{array}{c}
[\Pi \text{ is correct}] \\
\vdots \\
\{P_{\mathbf{p}_1}\} C_{\mathbf{p}_1} \{Q_{\mathbf{p}_1}\} \quad \dots \quad \{P_{\mathbf{p}_n}\} C_{\mathbf{p}_n} \{Q_{\mathbf{p}_n}\}
\end{array}
}{
\Pi \text{ is correct}
}
\\
\\
\text{(adapt)} \\
\frac{
\Pi \text{ is correct}
}{
\{P\} \text{ call } \mathbf{p}_i \{Q\}
} \quad \text{if } \models P \rightarrow P_{\mathbf{p}_i} \text{ and } \models Q_{\mathbf{p}_i} \rightarrow Q
\end{array}$$

From the point of view of verification conditions, the validity of those concerning a particular procedure \mathbf{p} is not sufficient to ensure the correctness of \mathbf{p} , since this relies on the correctness of the procedures invoked by \mathbf{p} . Instead, the verification conditions of a multi-procedure program should be considered as a whole, that is, as the union of the sets of VCs generated for each individual procedure, guaranteeing that each procedure keeps to its part of the contract, when invoked.

$$\text{VCG}(\Pi) = \bigcup_{i \in \{1, \dots, n\}} \text{VCG}(\{P_{\mathbf{p}_i}\} C_{\mathbf{p}_i} \{Q_{\mathbf{p}_i}\})$$

The following clauses complement accordingly the definition of the auxiliary functions of VCG:

$$\begin{aligned}
\text{wprec}(\text{call } \mathbf{p}_i, Q) &= P_{\mathbf{p}_i} \\
\text{VC}_{\text{aux}}(\text{call } \mathbf{p}_i, Q) &= \{[Q_{\mathbf{p}_i} \rightarrow Q]\}
\end{aligned}$$

Homeier and Martin's mechanically verified VCGen has been extended to a language with recursive procedures (with parameters) following a similar approach [48].

We end the section with some remarks on further aspects related to procedures.

Auxiliary Variables and Adaptation Rules. In the rules and verification conditions considered above for procedure calls, we admitted that contracts contained no occurrences of auxiliary variables. But this is a big limitation, since it precludes any kind of specification that relates the post-state and the pre-state, i.e. the outputs and inputs of procedures.

Consider again our *(adapt)* rule. If $Q_{\mathbf{p}_i}$ contains auxiliary variables, the side condition $Q_{\mathbf{p}_i} \rightarrow Q$ may be invalid for lack of information. Suppose $P_{\mathbf{p}}$ is $x == x_0$ and $Q_{\mathbf{p}}$ is $x == x_0 + 10$, P is $x > 0$ and Q is $x > 10$. Intuitively the triple $\{P\} \text{ call } \mathbf{p} \{Q\}$ is valid, but clearly $\not\models x == x_0 + 10 \rightarrow x > 10$.

The problem is that it is not possible to prove this implication between the contract's postcondition and the actual required postcondition, without taking into account information from P and $P_{\mathbf{p}}$: the two verification conditions should be replaced by a single one, combining two assertions that are interpreted in different states (the pre-state and the post-state of the procedure call), thus allowing information to be transported from one state to another. One way to solve this is via a syntactic encoding of one of the states, using variable substitutions and quantification. The following rule is adapted from [59].

$$\frac{\Pi \text{ is correct}}{\{P\} \text{ call } \mathbf{p}_i \{P\}} \quad \text{if } P \rightarrow \text{Forall } \vec{x}_f. (\text{Forall } \vec{y}_f. P_{\mathbf{p}}[\vec{y} \mapsto \vec{y}_f] \rightarrow Q_{\mathbf{p}}[\vec{y} \mapsto \vec{y}_f, \vec{x} \mapsto \vec{x}_f]) \rightarrow Q[\vec{x} \mapsto \vec{x}_f]$$

where \vec{y} are the auxiliary variables occurring in $\{P_{\mathbf{p}}\} C_{\mathbf{p}} \{Q_{\mathbf{p}}\}$, \vec{x} are the program variables of \mathbf{p} , and \vec{x}_f and \vec{y}_f are fresh variables.

The interaction between auxiliary variables and recursive procedures has long been a topic of investigation. Some of the first proposed solutions involved the use of modified structural rules. Rather than covering here in detail the theoretical issues involved (in particular how the properties of Hoare logic are affected by the introduction of an adaptation rule) and the solutions proposed over the years, we direct the reader to the surveys [2] and [21].

Kleymann [59] also offers a more recent study of the topic, where it is stressed that (i) the fundamental difficulty is the fact that auxiliary variables are not given a special status in Hoare logic; and (ii) the problem of adaptation is not specific of recursive procedures, and already manifests itself whenever one tries to apply the consequence rule of Hoare logic with a hypothesis containing auxiliary variables. Von Oheimb [82] explicitly addresses the treatment of mutual recursion in the presence of auxiliary variables.

Procedures with Parameters. We will not consider in detail the case of procedures with parameters; in the following we simply mention some of the issues that must be taken into account when generating verification conditions regarding the invocation of such procedures. Consider the following procedure definition, where x is the only formal parameter of \mathbf{p} , of type **int**.

$$\mathbf{proc } \mathbf{p}(x) = \{P_{\mathbf{p}}\} C_{\mathbf{p}} \{Q_{\mathbf{p}}\}$$

One novelty is that the variable x is *local* to the procedure (one could also have variables that are local to a *block* of code, but in this paper we have not considered this possibility). It must be possible for x to be used in the annotations, in particular in $P_{\mathbf{p}}$ or $Q_{\mathbf{p}}$, but also in invariants occurring inside $C_{\mathbf{p}}$, allowing them to refer to the value of the parameter.

Parameters have the effect of hiding global variables of the same name (both

in the code and in annotations), but this is not a big issue since they can be freely renamed. Renaming is essential when reasoning about procedures with parameters; it suffices to understand that when some other procedure \mathbf{q} executes the call $\mathbf{call\ p}(e)$, the expression e may contain occurrences of, say, a local variable y of \mathbf{q} that happens also to be a global variable occurring in $P_{\mathbf{p}}$. A verification condition will be generated containing occurrences of the same variable that have as origin different program variables.

Another issue that must be taken into account is that parameters passed by value should be treated differently from parameters passed by reference. In particular, occurrences of a parameter x passed by value in $Q_{\mathbf{p}}$, the procedure's postcondition, are usually interpreted as referring to the value of x in the pre-state in which the procedure was invoked, even if the programming language allows assigning locally the value of such parameters. If x is passed by reference on the other hand, occurrences of x in $Q_{\mathbf{p}}$ should be interpreted in the post-state, and an auxiliary variable (or an operator like `old` discussed before) should be used if it is necessary to refer to its value in the pre-state.

The most difficult issue created by the presence of parameters passed by reference is yet another manifestation of aliasing. If a procedure \mathbf{p} has two parameters x and y passed by reference, the invocation $\mathbf{p}(z, z)$, with z a global variable, creates aliasing since occurrences of both x and y in \mathbf{p} will refer to z , thus an assignment to x will modify also the value of y . In fact it is not even required that \mathbf{p} has two such parameters, since the global variable z may be used directly in the body of \mathbf{p} . This form of aliasing is known as *parameter aliasing*. In the C programming language parameters are passed by reference through pointers, thus parameter aliasing is reduced to pointer aliasing. Even when this is not the case, the techniques employed to reason about pointer aliasing can in general also be used for parameter aliasing.

Frame Properties. In languages like C or Java, evaluation of expressions may have *side effects* that modify the state of the program; `x++` is a typical example of such an expression. This issue has been addressed in most attempts to construct Hoare logics for realistic languages, see Section 11.

A procedure call is also an example of an expression with side effects. A command like $\mathbf{call\ p}$ alters the state of the program since execution of the body of \mathbf{p} is free to access global variables (or class and instance variables, in the case of object-oriented programs). To allow for modular reasoning, the procedure's contract can include information describing the unchanged part of the state, using a technique similar to that described at the end of Section 7. However, it is easy to see that this approach is inadequate. For instance, if new variables are introduced in the state, the specifications of many different procedures may have to be changed. This is usually known as

the *frame problem* [13], and it applies to all forms of subroutines, including procedures, functions, and object methods – it is in fact particularly important in object-oriented languages.

In general terms, it is more useful to include in contracts a description of what part of the state the subroutines are allowed to affect, rather than what part must remain unchanged. This is usually known as the subroutine’s *frame condition*. For instance in JML there are two frame annotations that may be included in specifications: **assignable** annotations, which list the variables that may be assigned by a method, and **modifies** annotations, which list variables whose value is allowed to be modified. Thus a variable that may be assigned inside a method, but whose value upon exit must be guaranteed to be restored to its entry value can be listed as assignable but not as modifiable.

Frame conditions are part of a method’s contract: typically they give rise to proof obligations when the method is being verified, and generate hypotheses that may be used when reasoning about calls to that method. In JML, methods that do not affect the state at all are called *pure*. A method annotated as pure can be used for specification purposes and called inside annotations.

11 Conclusion

We finish the paper with a brief survey of work that has contributed to bringing the benefits of deductive verification into practice.

Theorem Proving. As already mentioned, one reason for the reemergence of program verification is the tremendous progress that automated theorem provers have experienced in the last decade. Notably a new class of provers has appeared, called *Satisfiability Modulo Theories* (SMT) solvers [25, 26].

SMT solvers have their origins in *Boolean satisfiability* (SAT) solvers [67, 34]. While SAT solvers check efficiently the satisfiability of propositional formulas, SMT solvers check the satisfiability of first-order formulas containing operations from various theories such as the Booleans, bit-vectors, arithmetic, arrays, and recursive data types. More precisely, SMT solvers address the issue of satisfiability of quantifier-free first-order formulas (in conjunctive normal form), using as building blocks (i) a propositional SAT solver, and (ii) state-of-the-art theory solvers to study the satisfiability of sets of first-order literals.

SMT solvers have been traditionally used to support deductive software verification, and play an important role in modern verification architectures for programs annotated with contracts. They have also been applied with success

in several domains in computer science, such as model checking and automated test generation.

SMT solvers are the subject of very active research, and substantial technical advances have been taking place, propelled by the SMT-LIB initiative³ and the associated SMT-COMP⁴ competition, which provide a common input format and benchmarking framework for the evaluation of these systems. Popular solvers include Yices [31], CVC3 [7], Z3 [24], and Alt-Ergo [18].

In the absence of decidability results (recall that first-order logic is only semi-decidable) the importance of the user-guided deductive method cannot be overstated. Automated methods may never replace human genius in every situation, and some proofs appeal to the creativity of the user. *Interactive proof assistants* are tools aimed precisely at assisting users in the construction of derivations in a given proof system. Typically they can deal with very expressive logics, such as higher-order logic, and basically follow one of two approaches: an *axiomatic* approach, where the users are given the possibility to define the logic in which they desire to express proofs, and of which the Isabelle system [72] is a representative; or an *integrated* approach, where the basic language is sufficiently expressive to formulate most of mathematics, and of which the Coq system [12] is a representative.

Finally, we cite [79] as a recent survey on automated deduction for verification.

Automatic Generation of Invariants. The automated inference of invariants has been an active field of research since the 1970's. The first techniques proposed for inference of invariants were based on static methods, in particular abstract interpretation and a constraint-based approach [39, 57, 56, 80, 22, 20, 27]. More recent techniques also employ dynamic methods [33]. Loop invariants captured by these techniques can be extremely useful, but are typically insufficient to prove full functional correctness of programs, which usually requires human creativity.

Static techniques for invariant inference are used in modern static checkers such as Boogie, ESC and Why, see below.

Realistic Programming Languages. The gap between the point where this paper stops, and program verification for real-world languages begins, has been the object of significant work in recent years.

³ <http://www.smtlib.org/>

⁴ <http://www.smtcomp.org/>

One first aspect is that much of the recent work on program verification has been developed by a community federated around the *Java Modelling Language* (JML) [62, 60]. This is an annotation language for Java programs with support for preconditions, postconditions, frame conditions, and loop invariants, all written using the same syntax for expressions as the source language (exclusive of side-effects). JML also includes special keywords `\old`, with the interpretation explained in Section 7, and `\result`, which stands for the return value of a method.

The language incorporates many other modelling aspects that are useful to reason about object-oriented programs, such as class invariants (properties on the values of instance and class variables that must be preserved by all methods) and specification-only (or *model*) fields and methods, which allow for reasoning about hidden (or not yet implemented) specifications. It also includes a rich library of types with accompanying mathematical theories, including sets and sequences.

JML was not developed exclusively with program verification or other static checking methods in mind. In fact, the design-by-contract approach to software development is supported by a number of different tools for different tasks. These include dynamic assertion checkers, unit test class generators, and even generators that can be used to help write specifications. A number of program verification systems discussed below also use JML as specification language.

Admittedly, the material in Section 10 barely starts to cover the complex language constructs and memory models found in programming languages like C, Java or C#. For the particular case of object-oriented languages, there are many specific aspects that we have not covered at all. Poetzsch-Heffter and Müller [75] summarise as follows the difficulties involved in designing Hoare-style logics for these languages.

Three aspects make verification of OO-programs more complex than verification of programs with just recursive procedures and arbitrary pointer data structures: subtyping, abstract types, and dynamic binding. Subtyping allows variables to hold objects of different types. Abstract types have different or incomplete implementations. Thus, techniques are needed to formulate type properties without referring to implementations. Dynamic binding destroys the static connection between the calls and the body of a procedure.

A discussion of how these issues are addressed is beyond the scope of this paper; in the following we will concentrate instead on the organisational aspects. Accounts of the state of the art and current challenges in the field of object-oriented program verification (and specification languages) have been given by Jacobs, Kiniry and Warnier [51] and more recently by Leavens, Leino, and Müller [61].

The work on verification conditions for realistic languages falls roughly into two categories. Some researchers have proposed logics that attempt to fully describe the axiomatic semantics of the programming language; this is the case for instance of the following, which have been used to produce working program verification systems.

- Poetzsch-Heffter and Müller’s logic for sequential Java [75], used in the Jive [70] verification platform to generate verification conditions exported to the PVS [74] theorem prover. Following Gordon and Homeier and Martin, both the operational semantics of the language and the proposed Hoare inference system have been encoded in higher-order logic, and the latter proved sound with respect to the former.
- In the context of the Loop project, a Hoare logic for JML has been developed [52]. The approach followed in Loop is quite different from what has been discussed in this paper: a compiler first produces a logical theory (in the language of the PVS theorem prover) that describes the behaviour of a given JML-annotated Java program; program correctness is then proved in this theory. However, a weakest-precondition function can also be integrated in the tool [50], with the aim of reducing the required user interaction. This function is defined inside the logic, and works on translated programs.
- The KeY project’s JavaCardDL [10] is used in an integrated tool that supports the development of JML-annotated Java Card programs. JavaCardDL has not to our knowledge been mechanically proved correct with respect to a non-axiomatic semantics. This system uses a special-purpose theorem prover to construct proofs of JavaCardDL.

The architectures of the latter two systems do not use a separate VCGen component, and do not export proof obligations free of program constructs.

Von Oheimb’s Hoare logic for a subset of Java Card [83], although not used (to our knowledge) to produce a working tool, was an important development since it was mechanically proved sound and complete with respect to the operational semantics, using the Isabelle/HOL interactive theorem prover.

A second approach has concentrated on developing practical tools for program verification. The focus here has not been so much on correctness – these tools do not promise to find all errors a program may contain – but on usability, integration into the software development process, and compatibility with different proof tools.

We mention here two VCGen tools that are based on relatively small and simple *intermediate languages*, sufficiently expressive to support the translation of realistic languages. The more complex aspects are handled by an adequate translation into the intermediate language, rather than by the design of a complex Hoare logic for the object language. A good example of this is the memory

model: the intermediate language will typically have a very simple model with no heap, and the translation of the object language encodes the heap in some way into that model, to ensure an adequate treatment of references/pointers and aliasing.

The Boogie [5] program verifier was designed for the Spec# language (which stands for C# as JML stands for Java) and incorporates many advanced features, including integration with a development environment and automatic inference of loop invariants. The BoogiePL intermediate language is a guarded command language that features procedures and mutable variables, but excludes more complex features like methods, side-effects, or call-by-reference. So the translation of the object language into this intermediate language is responsible for encoding these features in a sound way.

The Why tool [36] was born from a very different perspective: the basic idea is to provide an interpretation of programs in type theory, such that the proof obligations for typability of the interpretation function coincide with the verification conditions. In this way, interpreting programs in type theory becomes an alternative to using a standard VCGen. The Why intermediate language is ML-like, including both functional and imperative features, exceptions, and labelling of statements (to make the use of auxiliary variables unnecessary). The type system accounts for *effects*, i.e. the type of an expression is annotated with, for instance, the set of variables that are possibly modified by its evaluation, which makes possible to eliminate aliasing. Why supports a variety of theorem provers, and has been used to produce program verification tools for both C (Caduceus [35], Frama-C [19]) and Java (Krakatoa [68]). Frama-C is a plugin-based general purpose static analyser for C programs, which works with the recently proposed *ANSI-C Specification Language* (ACSL) [9].

Finally, we remark that all the verification tools we have mentioned for realistic programming languages implement a form of *safety verification*. What we mean by this is that the tools generate some special verification conditions that are not directly related to the contracts included in the code. These conditions concern the safe execution of methods or procedures: by discharging them, one proves, for instance, that arithmetic errors (like division by zero) or inappropriate memory accesses do not take place. This is a comparatively easy form of verification, and the degree of automation achieved is usually high.

A treatment of safety at the theoretical level requires working with a semantics that deals with error states, and interpreting partial correctness as implying that a program does not terminate in an error state.

Acknowledgements

This work was supported by projects FAVAS (PTDC/EIA-CCO/105034/2008) and CROSS (PTDC/EIA-CCO/108995/2008), both funded by Fundação para a Ciência e Tecnologia (FCT).

References

- [1] Wolfgang Ahrendt, Thomas Baar, Bernhard Beckert, Richard Bubel, Martin Giese, Reiner Hähnle, Wolfram Menzel, Wojciech Mostowski, Andreas Roth, Steffen Schlager, and Peter H. Schmitt. The KeY tool. *Software and System Modeling*, 4(1):32–54, 2005.
- [2] Krzysztof R. Apt. Ten years of Hoare’s logic: A survey - part 1. *ACM Trans. Program. Lang. Syst.*, 3(4):431–483, 1981.
- [3] A. Azurat and I.S.W.B. Prasetya. A survey on embedding programming logics in a theorem prover. Technical report, University of Utrecht, 2002.
- [4] Roland Backhouse. *Program Construction – Calculating Implementations from Specifications*. John Wiley & Sons, Ltd, 2003.
- [5] Michael Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P. de Roever, editors, *FMCO*, volume 4111 of *Lecture Notes in Computer Science*, pages 364–387. Springer, 2005.
- [6] Michael Barnett and K. Rustan M. Leino. Weakest-precondition of unstructured programs. *SIGSOFT Softw. Eng. Notes*, 31(1):82–87, 2006.
- [7] Clark Barrett and Cesare Tinelli. CVC3. In Werner Damm and Holger Hermanns, editors, *Proceedings of the 19th International Conference on Computer Aided Verification (CAV ’07)*, volume 4590 of *Lecture Notes in Computer Science*, pages 298–302. Springer-Verlag, July 2007. Berlin, Germany.
- [8] Gilles Barthe, Pedro R. D’Argenio, and Tamara Rezk. Secure information flow by self-composition. In *CSFW*, pages 100–114. IEEE Computer Society, 2004.
- [9] Patrick Baudin, Pascal Cuoq, Jean-Christophe Filiâtre, Claude Marché, Benjamin Monate, Yannick Moy, and Virgile Prevosto. *ACSL: ANSI/ISO C Specification Language*. CEA LIST and INRIA, 2010.
- [10] Bernhard Beckert. A dynamic logic for the formal verification of Java Card programs. In Isabelle Attali and Thomas P. Jensen, editors, *Java Card Workshop*, volume 2041 of *Lecture Notes in Computer Science*, pages 6–24. Springer, 2000.

- [11] Yves Bertot. Theorem proving support in programming language semantics. *CoRR*, abs/0707.0926, 2007.
- [12] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. Springer Verlag, 2004.
- [13] Alexander Borgida, John Mylopoulos, and Raymond Reiter. On the frame problem in procedure specifications. *Software Engineering*, 21(10):785–798, 1995.
- [14] Richard Bornat. Proving pointer programs in Hoare logic. In Roland Carl Backhouse and José Nuno Oliveira, editors, *MPC*, volume 1837 of *Lecture Notes in Computer Science*, pages 102–126. Springer, 2000.
- [15] R. S. Boyer and J. S. Moore. *The Correctness Problem in Computer Science*, chapter A Verification Condition Generator for FORTRAN. Academic Press, 1981.
- [16] Rod Burstall. Some techniques for proving correctness of programs which alter data structures. *Machine Intelligence*, 7, 1972.
- [17] David R. Cok and Joseph Kiniry. ESC/Java2: Uniting ESC/Java and JML. In Gilles Barthe, Lilian Burdy, Marieke Huisman, Jean-Louis Lanet, and Traian Muntean, editors, *CASSIS*, volume 3362 of *Lecture Notes in Computer Science*, pages 108–128. Springer, 2004.
- [18] Sylvain Conchon, Evelyne Contejean, and Johannes Kanig. Ergo: a theorem prover for polymorphic first-order logic modulo theories, 2006.
- [19] Loc Correnson, Pascal Cuoq, Armand Puccetti, and Julien Signoles. Framac user manual. Available from the Framac website, <http://frama-c.com>, 2010.
- [20] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 84–97, Tucson, Arizona, 1978. ACM Press, New York, NY.
- [21] Patrick Cousot. Methods and logics for proving programs. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*, pages 841–994. Elsevier and MIT Press, 1990.
- [22] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM Symposium on Principles of Programming Languages*, pages 238–252, January 1977.
- [23] Ádám Darvas, Reiner Hähnle, and David Sands. A theorem proving approach to analysis of secure information flow. In Dieter Hutter and Markus Ullmann, editors, *SPC*, volume 3450 of *Lecture Notes in Computer Science*, pages 193–209. Springer, 2005.

- [24] Leonardo de Moura and Nikolaj Bjørner. *Z3: An Efficient SMT Solver*, volume 4963/2008 of *Lecture Notes in Computer Science*, pages 337–340. Springer Berlin, April 2008.
- [25] Leonardo de Moura and Nikolaj Bjrner. Satisfiability modulo theories: An appetizer. In Marcel Vinicius Medeiros Oliveira and Jim Woodcock, editors, *SBMF*, volume 5902 of *Lecture Notes in Computer Science*, pages 23–36. Springer, 2009.
- [26] Leonardo de Moura, Bruno Dutertre, and Natarajan Shankar. A tutorial on satisfiability modulo theories. In *Proceedings of the 19th international conference on Computer aided verification, CAV’07*, pages 20–36, Berlin, Heidelberg, 2007. Springer-Verlag.
- [27] Nachum Dershowitz and Zohar Manna. Inference rules for program annotation. In *Proceedings of the 3rd international conference on Software engineering, ICSE ’78*, pages 158–167, Piscataway, NJ, USA, 1978. IEEE Press.
- [28] David Detlefs, Greg Nelson, and James B. Saxe. Simplify: a theorem prover for program checking. *J. ACM*, 52(3):365–473, 2005.
- [29] E. W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, Englewood Cliffs, New Jersey, 1976.
- [30] Vijay D’Silva, Daniel Kroening, and Georg Weissenbacher. A survey of automated techniques for formal software verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 27(7):1165–1178, July 2008.
- [31] Bruno Dutertre and Leonardo De Moura. The Yices SMT solver. Technical report, SRI, 2006.
- [32] Matthew B. Dwyer, John Hatcliff, Robby Robby, Corina S. Pasareanu, and Willem Visser. Formal software analysis emerging trends in software model checking. In *2007 Future of Software Engineering, FOSE ’07*, pages 120–136, Washington, DC, USA, 2007. IEEE Computer Society.
- [33] Michael D. Ernst, Jake Cockrell, William G. Griswold, and David Notkin. Dynamically discovering likely program invariants to support program evolution. *IEEE Transactions on Software Engineering*, 27:99–123, 2001.
- [34] Niklas En and Niklas Srensson. An extensible SAT-solver. In Enrico Giunchiglia and Armando Tacchella, editors, *SAT*, volume 2919 of *Lecture Notes in Computer Science*, pages 502–518. Springer, 2003.
- [35] Jean-Christophe Filliâtre and Claude Marché. Multi-prover verification of C programs. In Jim Davies, Wolfram Schulte, and Michael Barnett, editors, *ICFEM*, volume 3308 of *Lecture Notes in Computer Science*, pages 15–29. Springer, 2004.
- [36] Jean-Christophe Filliâtre and Claude Marché. The Why/Krakatoa/Caduceus platform for deductive program verification. In Werner Damm and Holger

Hermanns, editors, *Proceedings of CAV'07*, volume 4590 of *Lecture Notes in Computer Science*, pages 173–177. Springer, 2007.

- [37] Cormac Flanagan and James B. Saxe. Avoiding exponential explosion: generating compact verification conditions. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 193–205, New York, NY, USA, 2001. ACM.
- [38] Robert Floyd. Assigning meaning to programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, number 19 in Proceedings of Symposia in Applied Mathematics, pages 19–32. American Mathematical Society, 1967.
- [39] Steven M. German and Ben Wegbreit. A synthesizer of inductive assertions. In *Proceedings of the May 19-22, 1975, national computer conference and exposition*, AFIPS '75, pages 369–376, New York, NY, USA, 1975. ACM.
- [40] Donald I. Good. Mechanical proofs about computer programs. Technical Report 41, The University of Texas at Austin, March 1984.
- [41] Michael J. C. Gordon. Mechanizing programming logics in higher order logic. In G. Birtwistle and P.A. Subrahmanyam, editors, *Current trends in hardware verification and automated theorem proving*, pages 387–439. Springer-Verlag New York, Inc., 1989.
- [42] David Gries. *The Science of Programming*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1987.
- [43] Reiner Hähnle and Richard Bubel. A Hoare-style calculus with explicit state updates. Department of Computer Science, Chalmers University of Technology.
- [44] D. Harel. Dynamic logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic Volume II — Extensions of Classical Logic*, pages 497–604. D. Reidel Publishing Company: Dordrecht, The Netherlands, 1984.
- [45] C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12:576–580, 1969.
- [46] C. A. R. Hoare. Procedures and parameters: An axiomatic approach. In *Proceedings of Symposium on Semantics of Algorithmic Languages*, volume 188 of *Lecture Notes in Mathematics*. Springer Berlin / Heidelberg, 1971.
- [47] Peter V. Homeier and David F. Martin. A mechanically verified verification condition generator. *Comput. J.*, 38(2):131–141, 1995.
- [48] Peter V. Homeier and David F. Martin. Mechanical verification of mutually recursive procedures. In Michael A. McRobbie and John K. Slaney, editors, *CADE*, volume 1104 of *Lecture Notes in Computer Science*, pages 201–215. Springer, 1996.
- [49] Shigeru Igarashi, Ralph L. London, and David C. Luckham. Automatic program verification I: A logical basis and its implementation. *Acta Inf.*, 4:145–182, 1974.

- [50] Bart Jacobs. Weakest pre-condition reasoning for Java programs with JML annotations. *J. Log. Algebr. Program.*, 58(1-2):61–88, 2004.
- [51] Bart Jacobs, Joseph Kiniry, and Martijn Warnier. Java program verification challenges. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P. de Roever, editors, *FMCO*, volume 2852 of *Lecture Notes in Computer Science*, pages 202–219. Springer, 2002.
- [52] Bart Jacobs and Erik Poll. A logic for the Java modeling language JML. In Heinrich Hußmann, editor, *FASE*, volume 2029 of *Lecture Notes in Computer Science*, pages 284–299. Springer, 2001.
- [53] Ranjit Jhala and Rupak Majumdar. Software model checking. *ACM Comput. Surv.*, 41:21:1–21:54, October 2009.
- [54] Cliff B. Jones. The early search for tractable ways of reasoning about programs. *IEEE Ann. Hist. Comput.*, 25(2):26–49, 2003.
- [55] Anne Kaldewaij. *Programming: the derivation of algorithms*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1990.
- [56] M. Karr. Affine relationships among variables of a program. *Acta Informatica*, 6:133–151, 1976.
- [57] Shmuel Katz and Zohar Manna. Logical analysis of programs. *Commun. ACM*, 19:188–206, April 1976.
- [58] James Cornelius King. *A program verifier*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA, 1969.
- [59] Thomas Kleymann. Hoare logic and auxiliary variables. *Formal Aspects of Computing*, 11(5):541–566, 1999.
- [60] G. Leavens and Y. Cheon. Design by Contract with JML, 2003.
- [61] Gary T. Leavens, K. Rustan M. Leino, and Peter Müller. Specification and verification challenges for sequential object-oriented programs. *Form. Asp. Comput.*, 19(2):159–189, 2007.
- [62] Gary T. Leavens, Clyde Ruby, K. Rustan, M. Leino, Erik Poll, and Bart Jacobs. JML: notations and tools supporting detailed design in Java. In *OOPSLA '00: Addendum to the 2000 proceedings of the conference on Object-oriented programming, systems, languages, and applications (Addendum)*, pages 105–106, New York, NY, USA, 2000. ACM.
- [63] K. Rustan M. Leino. Extended static checking: A ten-year perspective. In *Informatics - 10 Years Back. 10 Years Ahead.*, pages 157–175, London, UK, 2001. Springer-Verlag.
- [64] K. Rustan M. Leino. Efficient weakest preconditions. *Inf. Process. Lett.*, 93(6):281–288, 2005.

- [65] K. Rustan M. Leino, James B. Saxe, and Raymie Stata. Checking Java programs via guarded commands. In Ana M. D. Moreira and Serge Demeyer, editors, *Proceedings of ECOOP Workshops'99*, volume 1743 of *Lecture Notes in Computer Science*, pages 110–111. Springer, 1999.
- [66] Jacques Loeckx and Kurt Sieber. *The foundations of program verification (2nd ed.)*. John Wiley & Sons, Inc., New York, NY, USA, 1987.
- [67] Sharad Malik, Ying Zhao, Conor F. Madigan, Lintao Zhang, and Matthew W. Moskewicz. Chaff: Engineering an efficient SAT solver. *Design Automation Conference*, 0:530–535, 2001.
- [68] Claude Marché, Christine Paulin-Mohring, and Xavier Urbain. The KRAKATOA tool for certification of JAVA/JAVACARD programs annotated in JML. *J. Log. Algebr. Program.*, 58(1-2):89–106, 2004.
- [69] Bertrand Meyer. Applying “Design by Contract”. *IEEE Computer*, 25(10), 1992.
- [70] J. Meyer, P. Müller, and A. Poetzsch-Heffter. The JIVE system—implementation description. FernUniversität Hagen, 2000.
- [71] G. C. Necula. Proof-carrying code. In *Proceedings of POPL'97*, pages 106–119. ACM Press, 1997.
- [72] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [73] Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. Local reasoning about programs that alter data structures. In *Proceedings of the 15th International Workshop on Computer Science Logic, CSL '01*, pages 1–19, London, UK, 2001. Springer-Verlag.
- [74] Sam Owre, John M. Rushby, and Natarajan Shankar. PVS: A prototype verification system. In Deepak Kapur, editor, *CADE*, volume 607 of *Lecture Notes in Computer Science*, pages 748–752. Springer, 1992.
- [75] Arnd Poetzsch-Heffter and Peter Müller. A programming logic for sequential Java. In S. Doaitse Swierstra, editor, *ESOP*, volume 1576 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1999.
- [76] John C. Reynolds. *Theories of Programming Languages*. Cambridge University Press, Cambridge, England, 1998.
- [77] John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, pages 55–74. IEEE Computer Society, 2002.
- [78] A. Sabelfeld and A. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communications*, 21(1), 2003.
- [79] Natarajan Shankar. Automated deduction for verification. *ACM Comput. Surv.*, 41:20:1–20:56, October 2009.

- [80] Norihisa Suzuki and Kiyoshi Ishihata. Implementation of an array bound checker. In *Proceedings of the 4th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, POPL '77, pages 132–143, New York, NY, USA, 1977. ACM.
- [81] R. D. Tennent. *Specifying Software*. Cambridge University Press, New York, NY, USA, 2001.
- [82] David von Oheimb. Hoare logic for mutual recursion and local variables. In *Foundations of Software Technology and Theoretical Computer Science*, pages 168–180, 1999.
- [83] David von Oheimb. Hoare logic for Java in Isabelle/HOL. *Concurrency and Computation: Practice and Experience*, 13(13):1173–1214, 2001.
- [84] Glynn Winskel. *The Formal Semantics of Programming Languages: An introduction*. Foundations of Computing. The MIT Press, Cambridge, Massachusetts, 1993.