

Formally Validating a Practical Verification Condition Generator

Gaurav Parthasarathy¹, Peter Müller¹, and Alexander J. Summers²

¹ Department of Computer Science, ETH Zurich, Switzerland
{gaurav.parthasarathy,peter.mueller}@inf.ethz.ch

² University of British Columbia
alex.summers@ubc.ca



Abstract. A program verifier produces reliable results only if both the *logic* used to justify the program’s correctness is sound, and the *implementation* of the program verifier is itself correct. Whereas it is common to formally prove soundness of the logic, the implementation of a verifier typically remains unverified. Bugs in verifier implementations may compromise the trustworthiness of successful verification results. Since program verifiers used in practice are complex, evolving software systems, it is generally not feasible to formally verify their implementation. In this paper, we present an alternative approach: we *validate successful runs* of the widely-used Boogie verifier by producing a *certificate* which proves correctness of the obtained verification result. Boogie performs a complex series of program translations before ultimately generating a verification condition whose validity should imply the correctness of the input program. We show how to certify three of Boogie’s core transformation phases: the elimination of cyclic control flow paths, the (SSA-like) replacement of assignments by assumptions using fresh variables (passification), and the final generation of verification conditions. Similar translations are employed by other verifiers. Our implementation produces certificates in Isabelle, based on a novel formalisation of the Boogie language.

1 Introduction

Program verifiers are tools which attempt to prove the correctness of an implementation with respect to its specification. A successful verification attempt is, however, only meaningful if both the *logic* used to justify the program’s correctness is sound, and the *implementation* of the program verifier is itself correct. It is common to formally prove soundness of the logic, but the implementations of program verifiers typically remain unverified. As is standard for complex software systems, bugs in verifier implementations can and do arise, potentially raising doubts as to the trustworthiness of successful verification results.

One way to close this gap is to prove a verifier’s implementation correct. However, such a *once-and-for-all* approach faces serious challenges. Verifying an existing implementation bottom-up is not practically feasible because such implementations tend to be large and complex (for instance, the Boogie verifier [29] consists of over 30K lines of imperative C# code), use a variety of

libraries, and are typically written in efficient mainstream programming languages which themselves lack a formalisation. Alternatively, one could develop a verifier that is correct by construction. However, this approach requires the verifier to be (re-)implemented in an interactive theorem prover (ITP) such as Coq [14] or Isabelle [24]. This precludes the free choice of implementation language and paradigm, exploitation of concurrency, and possibility of tight integration with standard compilers and IDEs, which is often desirable for program verifiers [4, 5, 13, 26]. Both verification approaches substantially impede software maintenance, which is problematic since verifiers are often rapidly-evolving software projects (for instance, the Boogie repository [1] contains more than 5000 commits).

To address these challenges, in this work we employ a different approach. Instead of verifying the implementation once and for all, we *validate specific runs* of the verifier by automatically producing a *certificate* which proves the correctness of the obtained verification result. Our certificate generation formally relates the input and output of the verifier, but does so largely independently of its implementation, which can freely employ complex languages, algorithms, or optimisations. Our certificates are formal proofs in Isabelle, and so checkable by an independent trusted tool; their guarantees for a certified run of the verifier are as strong as those provided by a (hypothetical) verified verifier.

We apply our novel verifier validation approach to the widely-used Boogie verifier, which verifies programs written in the intermediate verification language Boogie. The Boogie verifier is a *verification condition generator*: it verifies programs by generating a verification condition (VC), whose validity is then discharged by an SMT solver. Certifying a verifier run requires proving that validity of the VC implies the correctness of the input program. Certification of the validity-checking of the VC is an orthogonal concern; our results can be combined with work in that area [11, 15, 19] to obtain end-to-end guarantees.

Like many automatic verifiers, Boogie is a *translational verifier*: it performs a sequence of substantial Boogie-to-Boogie translations (*phases*), simplifying the task and output of the final efficient VC computation [6, 18]. The key challenges in certifying runs of the Boogie tool are to certify each of these phases, including final VC generation. In particular, we present novel techniques for making the following three key phases (and many smaller ones) of Boogie’s tool chain certifying:

1. The elimination of loops (more precisely, cycles in the CFG) by reducing the correctness of loops to checking loop invariants (*CFG-to-DAG phase*)
2. The replacement of assignments by (SSA-style) introduction of fresh variables and suitable **assume** statements (*passification phase*)
3. The final generation of the VC, which includes the erasure and logical encoding of Boogie’s polymorphic type system [33] (*VC phase*).

The certification of such verifier phases is related to existing work on compiler verification [34] and validation [8, 41, 42]. However, the translations and the certified property we tackle here are fundamentally different from those in compilers. Compilers typically require that each execution of the target program

corresponds to an execution of the source program. In contrast, the encoding of a program in a translational verifier typically has intentionally more executions (for instance, allows more non-determinism). Moreover, translational verifiers need to handle features not present in standard programming languages such as **assume** statements and background theories. Prior work on validating such verifier phases has been limited in the supported language and extent of the formal guarantee; we discuss comparisons in detail in Sec. 8.

Contributions. Our paper makes the following technical contributions.

1. The first formal semantics for a significant subset of Boogie (including axioms, polymorphism, type constructors), mechanised in Isabelle.
2. A validation technique for two core program-to-program translations occurring in verifiers (CFG-to-DAG and passification).
3. A validation technique for the VC phase, handling polymorphism erasure and Boogie’s type system encoding [31], for which no prior formal proof exists.
4. A version of the Boogie implementation that produces certificates for a significant subset of Boogie.

Making the Boogie verifier certifying is an important result, reducing the trusted code base for a wide variety of verification tools implemented via encodings into Boogie, e.g. Dafny [31], VCC [13], Corral [28], and Viper [35]. Moreover, the technical approach we present here can in future be applied to the certification of the translations performed by these tools, and those based on comparable intermediate verification languages such as Frama-C [26] and Krakatoa [17] based on Why3 [16] and Prusti [4] and VerCors [10] based on Viper [35].

Outline. Sec. 2 explains at a high-level, how our validation approach is structured for the different phases. Sec. 3 introduces a formal semantics for Boogie. Secs. 4, 5 and 6 present our validation of the CFG-to-DAG, passification, and VC phases, respectively. Sec. 7 evaluates our certificate-producing version of Boogie. Sec. 8 discusses related work. Sec. 9 concludes. Further details are available in our accompanying technical report (hereafter, TR) [37].

2 Approach

A Boogie program consists of a set of procedures, each with a specification and a procedure body in the form of a (reducible) control-flow-graph (CFG), whose blocks contain basic commands; we present the formal details in the next section. Boogie verifies each procedure modularly, desugaring procedure calls according to their specifications. Verification is implemented via a series of phases: program-to-program translations and a final computation of a VC to be checked by an SMT solver. Our goal is to formally certify (per run of Boogie) that validity of this VC implies the correctness of the original procedure.

To keep the complexity of certificates manageable, our technical approach is *modular* in three dimensions: decomposing our formal goal per *procedure* in the




Fig. 1. Key phases of verification in Boogie and their certification. The solid edges show Boogie’s transformations on a procedure body; each node G_i represents a control-flow graph. Our final certificate (dashed edge) is constructed by formally linking the three phase certificates represented by the dotted edges. Each of the three phase certificates also incorporate extra smaller transformations that we do not show here.

Boogie program, per *phase* of the Boogie verification, and per *block* in the CFG of each procedure. This modularity makes the full automation of our certification proofs in Isabelle practical. In the following, we give a high-level overview of this modular structure; the details are presented in subsequent sections.

Procedure decomposition. Boogie has no notion of a main program or an overall program execution. A Boogie program is correct if each of its procedures is individually correct (that is, the procedure body has no failing traces, as we make precise in the next section). Boogie computes a separate VC for each procedure, and we correspondingly validate the verification of each procedure separately.

Phase decomposition. We break our overall validation efforts down into per-phase sub-problems. In this paper, we focus on the following three most substantial and technically-challenging of these sequential phases, illustrated in Fig. 1. (1) The *CFG-to-DAG phase* translates a (possibly-cyclic) CFG to an acyclic CFG (*cf.* Sec. 4). This phase substantially alters the CFG structure, cutting loops using annotated loop invariants to over-approximate their executions. (2) The *passification phase* eliminates imperative updates by transforming the code into static single assignment (SSA) form and then replacing assignments with *constraints* on variable versions (*cf.* Sec. 5). Both of these phases introduce extra non-determinism and **assume** statements (which, if implemented incorrectly could make verification unsound by masking errors in the program). (3) The final *VC phase* translates the acyclic, passified CFG to a verification condition that, in addition to capturing the weakest precondition, encodes away Boogie’s polymorphic type system [33].

We construct certificates for each of these key phases separately (depicted by the blue dotted lines in Fig. 1). For each phase, we certify that *if* the target of the translation phase is correct (a correct Boogie program for the first two phases; a valid VC for the VC phase) then the source (program) of the phase is correct. This modular approach lets us focus the proof strategy for each phase on its conceptually-relevant concerns, and provides robustness against *changes* to the verifier since at most the certification of the changed phases may need adjustment. Logically, our per-phase certificates are finally glued together to guarantee the analogous end-to-end property for the entire pipeline, depicted by

the green dashed edge in Fig. 1. For our certificates, we import the input and output programs (and VC) of each key phase from Boogie into Isabelle; we do not reimplement any of Boogie’s phases inside Isabelle.

The certificates of the key phases also incorporate various smaller transformations between the key phases, such as peephole optimisation. Our work also validates these smaller transformations, but we focus the presentation on the key phases in this paper. Boogie also performs several smaller translation steps *prior* to the CFG-to-DAG phase. These include transforming ASTs to corresponding CFGs, optimisations such as dead variable elimination, and desugaring procedure calls using their specifications (via explicit **assert**, **assume**, and **havoc** statements). Our approach applies analogously to these initial smaller phases, but our current implementation certifies only the pipeline of all phases from the (input to the) CFG-to-DAG phase onwards. Thus, our certificate relates Boogie’s VC to the original source AST program so long as these prior translation steps are correct.

CFG decomposition. When tackling the certification of *each* phase, we further break down validation of a procedure’s CFG in the source program of the phase into sub-problems for each block in the CFG. We prove two results for each block in the source CFG:

1. *Local block lemmas:* We prove an independent lemma for each source CFG block in isolation, relating the executions through the block with the corresponding block in the target program (or the VC generated for that block, in the case of the VC phase). In particular, this lemma implies that if the target block has no failing executions (or the VC generated for that block holds, for the VC phase), neither does the source block for corresponding input states.
2. *Global block theorems:* We show analogous per-block results concerning all executions *from this block onwards* extending to the end of the procedure in question; we build these compositionally by reverse-topological traversal of either the source or target CFGs, as appropriate. The global block theorem for the entry block establishes correctness of the phase.

This decomposition separates command-level reasoning (local block lemmas) from CFG-level reasoning (global block theorems). It enables concise lemmas and proofs in Isabelle and makes each comprehensible to a human.

3 A Formal Semantics for Boogie

Our certificates prove that the validity of a VC generated by Boogie formally implies correctness of the Boogie CFG-to-DAG source program. This proof relies crucially on a formal semantics for Boogie itself. Our first contribution is the first such formal semantics for a significant subset of Boogie, mechanised in Isabelle. Our semantics uses the Boogie reference manual [29], the presentation of its type system [33], and the Boogie implementation for reference; none of those provide a formal account of the language. For space reasons, we explain only the key concepts of our detailed formalisation here; more details are provided in App.

$$\begin{aligned}
e ::= & x \mid \mathbf{false} \mid \mathbf{true} \mid i \mid e_1 \mathit{bop} e_2 \mid \mathit{uop}(e) \mid f[\vec{\tau}](\vec{e}) \mid \mathbf{old}(e) \mid \\
& \forall x : \tau. e \mid \exists x : \tau. e \mid \forall_{ty} t. e \mid \exists_{ty} \tau. e \\
\tau ::= & \mathit{Int} \mid \mathit{Bool} \mid C(\vec{\tau}) \mid t \quad c ::= \mathbf{assume} e \mid \mathbf{assert} e \mid x := e \mid \mathbf{havoc} x
\end{aligned}$$

Fig. 2. The syntax of our formalised Boogie subset, where τ , e , and c , denote the types, expressions, and basic commands respectively; control-flow is handled via CFGs over the basic commands. bop and uop denote binary and unary operations, respectively.

A of the TR [37] and the full Isabelle mechanisation is available as part of our accompanying artifact [36].

3.1 The Boogie Language

Boogie programs consist of a set of top-level declarations of global variables and constants (the *global data*), axioms, uninterpreted (polymorphic) functions, type constructors, and procedures. A procedure declaration includes parameter, local-variable, and result-variable declarations (the *local data*), a pre- and postcondition, and a procedure body given as a CFG.³ CFGs are formalised as usual in terms of basic blocks (containing a possibly-empty list of *basic commands*), and edges; semantically, execution after a basic block continues via any of its successors non-deterministically.

The types, expressions, and basic commands in our Boogie subset are shown in Fig. 2. We support the primitive types *Int* and *Bool*; types obtained via declared type constructors are *uninterpreted types*; the sets of values such types denote are constrained only via Boogie axioms and **assume** commands. Moreover, types can contain type variables (for instance, to specify polymorphic functions).

Boogie expression syntax is largely standard (e.g. including typical arithmetic and boolean operations). Old-expressions **old**(e) evaluate the expression e w.r.t. the current local data and the global data as it *was* in the pre-state of the procedure execution. Boogie expressions also include universal and existential *value* quantification (written $\forall x : \tau. e$ and $\exists x : \tau. e$), as well as universal and existential *type* quantification (written $\forall_{ty} t. e$ and $\exists_{ty} \tau. e$). In the latter, t is bound in e and quantifies over *closed* Boogie types (i.e. types that do not contain any type variables).

Basic commands form the single-steps of traces through a Boogie CFG; sequential composition is implicit in the list of basic commands in a CFG basic block and further control flow (including loops) is prescribed by CFG edges. Boogie’s basic commands are assumes, asserts, assignments, and havocs; **havoc** x non-deterministically assigns a value matching the type of variable x to x .

³ Source-level procedure specifications also include *modifies clauses*, declaring a set of global variables the procedure may modify. As we tackle Boogie programs after procedure calls have been desugared, there are no *modifies clauses* in our formalisation.

The main Boogie features *not* supported by our subset are maps and other primitive types such as bitvectors. Boogie maps are polymorphic and impredicative, i.e. one can define maps that contain themselves in their domain. Giving a semantic model for such maps in a proof assistant such as Isabelle or Coq is non-trivial; we aim to tackle this issue in the future. Modelling bitvectors will be simpler, although maintaining full automation may require some additional work.

3.2 Operational Semantics

Values and state model. Our formalisation embeds integer and boolean values shallowly as their Isabelle counterparts; an Isabelle carrier type for all *abstract values* (those of uninterpreted types) is a parameter of our formalisation. Each uninterpreted type is (indirectly) associated with a *non-empty* subset of abstract values via a *type interpretation* map \mathcal{T} from abstract values to (single) types; particular interpretations of uninterpreted types can be obtained via different choices of type interpretation \mathcal{T} .

One can understand Boogie programs in terms of the sets of possible *traces* through each procedure body. Traces are (as usual) composed of sequences of steps according to the semantics of basic commands and paths through the CFG; these can be finite or infinite (representing a non-terminating execution). A trace may halt in three cases: (1) an exit block of the procedure is reached in a state satisfying the procedure’s postcondition (a *complete* trace),⁴ (2) an **assert** A command is reached in a state not satisfying assertion A (a *failing* trace), or (3) an **assume** A command is reached in a state not satisfying A (a trace which *goes to magic* and stops). Our formalisation correspondingly includes three kinds of Boogie program states: a distinguished *failure state* F , a distinguished *magic state* M , and *normal states* $N((os, gs, ls))$. A normal state is a triple of partial mappings from variables to values for the old global state (for the evaluation of old-expressions), the (current) global state, and the local state, respectively.

Expression evaluation. An expression e evaluates to value v if the (big-step) judgement $\mathcal{T}, A, \Gamma, \Omega \vdash \langle e, N(ns) \rangle \Downarrow v$ holds in the context $(\mathcal{T}, A, \Gamma, \Omega)$. Here, \mathcal{T} is a *type interpretation* (as above), A is a *variable context*: a pair (G, L) of type declarations for the global (G) and local (L) data. Γ is a *function interpretation*, which maps each function name to a semantic function mapping a list of types and a list of values to a return value. The type substitution Ω maps type variables to types.

The rules defining this judgement can be found in App. A.2 of the TR [37]. For example, the following rule expresses when a universal type quantification evaluates to **true** (t is bound to the quantified type and may occur in e):

$$\frac{\forall \tau. \text{closed}(\tau) \implies \mathcal{T}, A, \Gamma, \Omega(t \mapsto \tau) \vdash \langle e, ns \rangle \Downarrow \mathbf{true}}{\mathcal{T}, A, \Gamma, \Omega \vdash \langle \forall_{ty} t. e, ns \rangle \Downarrow \mathbf{true}}$$

⁴ The case of the postcondition *not* holding is subsumed under point (2), since Boogie checks postconditions by generating extra **assert** statements.

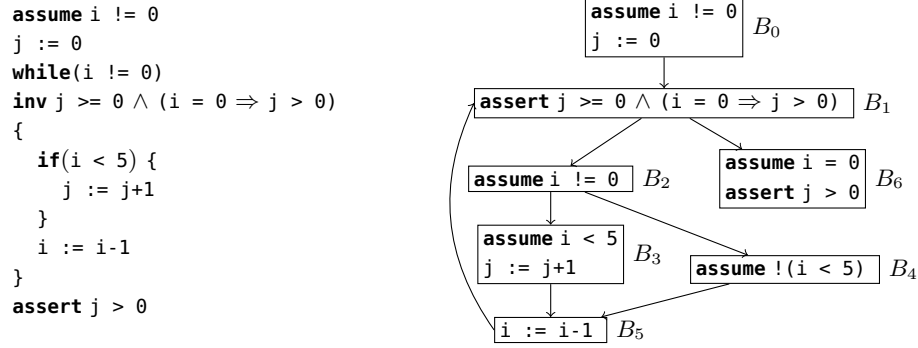


Fig. 3. Running example in source code and CFG representation, respectively.

The premise requires one to show that the expression e reduces to **true** for every possible type τ that is closed. In general, expression evaluation is possible only for well-typed expressions; we also formalise Boogie’s type system and (for the first time) prove its type safety for expressions in Isabelle.

Command and CFG reduction. The (big-step) judgement $\mathcal{T}, \Lambda, \Gamma, \Omega \vdash \langle c, s \rangle \rightarrow s'$ defines when a command c reduces in state s to state s' ; the rules are in App. A.3 of the TR [37]. This reduction is lifted to lists of commands cs to model the semantics of a single trace through a CFG block (the judgement $\mathcal{T}, \Lambda, \Gamma, \Omega \vdash \langle cs, s \rangle [\rightarrow] s'$). The operational semantics of CFGs is modelled by the (small-step) judgement $\mathcal{T}, \Lambda, \Gamma, \Omega, G \vdash \delta \rightarrow_{\text{CFG}} \delta'$, expressing that the CFG configuration δ reduces to configuration δ' in the CFG G . A CFG configuration is either *active* or *final*. An active configuration is given by a tuple $(\text{inl}(b_n), s)$, where b_n is the block identifier indicating the current position of the execution and s is the current state. A final configuration consists of a tuple $(\text{inr}(()), s)$ for state s (and unit value ()) and is reached at the end of a block that has either no successors, or is in a magic or failure state.

3.3 Correctness

A procedure is *correct* if it has *no failing traces*. This is a *partial correctness* semantics; a procedure body whose traces never leave a loop is trivially correct provided that no intermediate **assert** commands fail. Procedure correctness relies on CFG correctness. A CFG G is correct w.r.t. a postcondition Q and a context $(\mathcal{T}, \Lambda, \Gamma, \Omega)$ in an initial normal state $\mathbf{N}(ns)$ if the following holds for all configurations (r, s') :

$$\begin{aligned} \mathcal{T}, \Lambda, \Gamma, \Omega, G \vdash (\text{inl}(\text{entry}(G)), \mathbf{N}(ns)) \rightarrow_{\text{CFG}}^* (r, s') \implies [s' \neq \text{F} \wedge \\ (r = \text{inr}(())) \implies (\forall ns'. s' = \mathbf{N}(ns') \implies \mathcal{T}, \Lambda, \Gamma, \Omega \vdash \langle Q, \mathbf{N}(ns') \rangle \Downarrow \text{true})] \end{aligned}$$

where $\text{entry}(G)$ is the entry block of G and $\rightarrow_{\text{CFG}}^*$ is the reflexive-transitive closure of the CFG reduction. The postcondition is needed only if a final configuration

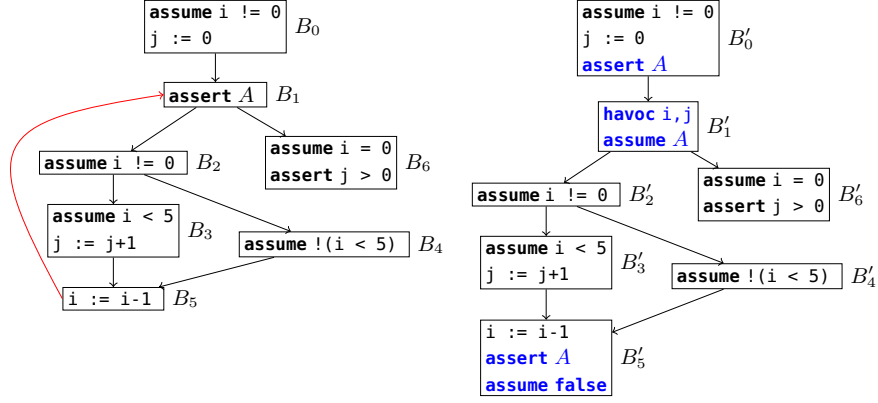


Fig. 4. The CFG-to-DAG phase applied to the running example (source is left, target is right). The back-edge (the red edge from B_5 to B_1 in the left CFG) is eliminated. The blue commands are new. A is given by $j \geq 0 \wedge (i = 0 \Rightarrow j > 0)$.

is reached in a normal state, while failing states must be unreachable. Whenever we omit Q , we implicitly mean the postcondition to be simply **true**. In our tool, we consider only empty initial mappings Ω , since we do not support procedure type parameters (lifting our work to this feature will be straightforward).

For a procedure p to be correct w.r.t. a context, its body CFG must be correct w.r.t. the same context and p 's postcondition, *for all* initial normal states $N(ns)$ that satisfy p 's precondition and which respect the context. For ns to *respect* a context, it must be well-typed and must satisfy the axioms when restricted to its constants. We say that p is *correct*, if it is correct w.r.t. *all well-formed contexts*, which must have a well-typed function interpretation and a type interpretation that inhabits every uninterpreted closed type (and only those).

Running example. We will use the simple CFG of Fig. 3 as a running example, intended as body of a procedure with trivial (**true**) pre- and post-conditions. The code includes a simple loop with a declared loop invariant, which functions as a classical Floyd/Hoare-style inductive invariant, and for the moment can be considered as an implicit **assert** statement at the loop head. The CFG has infinite traces: those which start from any state in which i is negative. Traces starting from a state in which i is zero go to magic; they do not reach the loop. The program is correct (has no failing traces): all other initial states will result in traces that satisfy the loop invariant and the final **assert** statement. If we removed the initial **assume** statement, however, there *would* be failing traces: the loop invariant check would fail if i were initially zero.

4 The CFG-to-DAG Phase

In this section, we present the validation for the CFG-to-DAG phase in the Boogie verifier. This phase is challenging as it changes the CFG structure, inserts

additional non-deterministic assignments and **assume** statements, and must do so correctly for arbitrary (reducible) nested loop structures, which can include unstructured control flow (e.g. jumps out of loops).

4.1 CFG-to-DAG Phase Overview

The CFG-to-DAG phase applies to every *loop head* block identified by Boogie’s implementation and any *back-edges* from a block reachable from the loop head block back to the loop head (following standard definitions for reducible CFGs [21]). Fig. 4 illustrates the phase’s effect on our running example. Block B_1 is the (only) loop head here, and the edge from B_5 to it the only back-edge (completing looping paths via B_2 and B_3 or B_2 and B_4). An **assert** A statement starting a loop head (like B_1) is interpreted as declaring A to be the loop invariant.⁵ The CFG-to-DAG phase performs the following steps:

1. Accumulate a set X_H of all (local and global) variables *assigned-to* on *any looping path* from the loop head back to itself. In our example, X_H is $\{i, j\}$.
2. Move the **assert** A statement declaring a loop invariant (if any) from the loop head to the end of *each preceding* block (in our example: B_0 and B_5).
3. Insert **havoc** statements at the start of the loop head block per variable in X_H , followed by a single **assume** A statement (preceding any further statements).
4. For each block with a back-edge to the loop head, delete the back-edge; if this leaves the block with no successors, append **assume false** to its commands.⁶

The havoc-then-assume sequence introduced in step 3 can be understood as generating traces for *arbitrary values* of X_H satisfying the loop invariant A , effectively over-approximating the set of states reachable at the loop head in the original program. In particular, the remnants of any originally looping path (e.g. B'_1, B'_2, B'_3, B'_5) enforce that any non-failing trace starting from any such state must (due to the **assert** added to block B'_5 in step 2) result in a state which re-establishes the loop invariant. Such paths exist only to enforce this inductive step (analogously to the premise of a Hoare logic while rule); so long as the **assert** succeeds, we can discard these traces via step 4.

While we illustrate this step on a simple CFG, in general a loop head may have multiple back-edges, looping structures may nest, and edges may exit multiple loops. For the above translation to be correct, the CFG must be reducible and loop heads and corresponding back-edges identified accurately, which is complex in general. Importantly (but perhaps surprisingly), our work makes this phase of Boogie certifying *without* explicitly verifying (or even defining) these notions.

4.2 CFG-to-DAG Certification: Local Block Lemmas

We define first our local block lemmas for this phase. Recall that these prove that if executing the statements of a target block yields no failing executions,

⁵ In general, multiple asserts at the beginning of a loop head may form the invariant.

⁶ Omitting **assume false** if there are no successors would be incomplete, since otherwise the postcondition would have to be satisfied.

the same holds for the corresponding source block; this result is trivial for source blocks other than loop heads and their immediate predecessors, since these are unchanged in this phase. To enable eventual composition of our block lemmas, we need to also reflect the role of the **assume** and **assert** statements employed in this phase. The formal statement of our local block lemmas is as follows⁷:

Theorem 1 (CFG-to-DAG Local Block Lemma). *Let B be a source block with commands cs_S , whose corresponding target block has commands cs_T . If B is a loop head, let X_H be as defined in CFG-to-DAG step 1 (and empty otherwise) and let A_{pre} be its loop invariant (or **true** otherwise). If B is a predecessor of a loop head, let A_{post} be the loop invariant of its successor (and **true** otherwise). Then, if:*

1. $\mathcal{T}, \Lambda, \Gamma, \Omega \vdash \langle cs_S, N(ns_1) \rangle [\rightarrow] s'_1$
2. $\forall s'_2. \mathcal{T}, \Lambda, \Gamma, \Omega \vdash \langle cs_T, N(ns_2) \rangle [\rightarrow] s'_2 \implies s'_2 \neq F$
3. A_{pre} is satisfied in ns_1 , and ns_2 differs from ns_1 only on variables in X_H and variables not defined in Λ

then: $s'_1 \neq F$ and if s'_1 is a normal state, then (1) A_{post} is satisfied in s'_1 , and (2) if no **assume false** was added at the end of cs_T , then there is a target execution in cs_T from $N(ns_2)$ that reaches a normal state that differs from s'_1 only on variables not defined in Λ .

The gist of this lemma is to capture *locally* the ideas behind the four steps of the phase. For example, consequence (1) reflects that *after* the transformation, any blocks that *were previously* predecessors of a loop head (B'_0 and B'_5 in our running example) will have an **assert** statement checking for the corresponding invariant (and so if the target program has no failing traces, in each trace this invariant will be true at that point).

4.3 CFG-to-DAG Certification: Global Block Theorems

We lift our certification to *all* traces through the source and target CFGs; the statement of the corresponding global block theorems is similar to that of local block theorems lifted to CFG executions, and for space reasons we do not present it here, but it is included in our Isabelle formalisation. In particular, we prove for each block (working in reverse topological order through the target CFG blocks) that if executions starting in the target CFG block never fail, neither do any executions starting from the corresponding source CFG block, and looping paths modify at most the variables havoced according to step 3 of the phase.

The major challenge in these proofs is reasoning about looping paths in the source CFG, since these revisit blocks. To solve this challenge, we perform inductive arguments per loop head in terms of the number of steps remaining in the trace in question.⁸ Our global block theorem for a block B then carries as

⁷ We omit some details regarding well-typedness, handled fully in our formalisation.

⁸ This may seem insufficient since traces can be infinite, but importantly a *failing* trace is always finite, and our theorems need only eliminate the chance of failing traces.




Fig. 5. The passification phase applied to the branch in the running example with the result on the right. The final (green) commands in B''_3 and B''_4 are the synchronisation commands. At the uppermost blocks shown here, the current versions of i and j are $i1$ and $j2$, respectively. The full CFGs are shown in App. B of the TR [37].

an assumption an induction hypothesis for each loop that contains B . Proving a global block theorem for the origin of a back-edge is taken care of by applying the corresponding induction hypothesis.

This proof strategy works only if we have obtained the induction hypothesis for the loop head *before* we use the global block theorem of the origin of a back-edge (otherwise we cannot discharge the block theorem’s hypothesis). In other words, our proof implicitly shows the necessary requirement that loop heads (as identified by Boogie) dominate all back-edges reaching them *without us formalising any notion of domination, CFG reducibility, or any other advanced graph-theoretic concept*. This shows a major benefit of our validation approach over a once-and-for-all verification of Boogie itself: our proofs indirectly check that the identification of loop heads and back-edges guarantees the necessary *semantic properties* without being concerned with *how* Boogie’s implementation computes this information.

Our approach applies equally to nested loops and more-generally to reducible CFG structures; *all* corresponding induction hypotheses are carried through from the visited loop heads. The requirement that no more than the haved variables X_H are modified in the source program is easily handled by showing that variables modified in an inner loop are a subset of those in outer loops. As for all of our results, our global block lemmas are proven automatically in Isabelle per Boogie procedure, providing per-run certificates for this phase.

5 The Passification Phase

In this section, we describe the validation of the passification phase in the Boogie verifier. Unlike the previous phase, passification makes no changes to the CFG structure, but makes substantial changes to the program states (via SSA-like renamings), substantially increases non-determinism, and employs **assume** statements to re-tame the sets of possible traces.

5.1 Passification Phase Overview

The main goal of passification is to eliminate assignments such that a more efficient VC can be ultimately generated [6, 18, 30]. In the Boogie verifier, this is implemented as a single transformation phase that can be thought of as two independent steps. Firstly, the source CFG is transformed into *static single assignment* (SSA) form, introducing *versions* (fresh variables) for each original program variable such that each version is assigned at most once in any program trace. In a second step, variable assignments are *completely eliminated*: each assignment command $x := e$ is replaced by **assume** $x = e$. Havoc statements are simply removed; their effect is implicit in the fact that a new variable version is used (via the SSA step) *after* such a statement.

Fig. 5 shows the effect of this phase on four blocks of our running example (the full figure of the target CFG is shown in App. B of the TR [37]). The commands inserted just before the join block (here, B_5'') introduce a consistent variable version (here, j4) for use in the join block. It is convenient to speak of target variables in terms of their source program counterparts: we say e.g. that j has *version 4* on entry to block B_5' .

Compared to traces through the source program, the space of variable values in a trace through the target program is initially much larger; each version may, on entry to the CFG, have an arbitrary value. For example, j4 may have any value on entry to B_2'' ; traces in which its value does not correspond to the constraint of the **assume** statements in B_3'' or B_4'' will go to magic and not reach B_5'' . Importantly, however, not *all* traces go to magic; enough are preserved to simulate the executions of the original program: each **assume** statement constrains the value of exactly one variable version, and the same version is never constrained more than once. Capturing this delicate argument formally is the main challenge in certifying this step.

As extra parts of the passification phase, the Boogie verifier performs constant propagation and desugars old-expressions (using variable versions appropriate to the entry point of the CFG). We omit their descriptions here for brevity, but our implementation certifies them.

5.2 Passification Certification: Local Block Lemmas

To validate the passification phase, it is sufficient to show that each source execution is simulated by a corresponding target execution, made precise by constructing a relation between the states in these executions. Such *forward simulation* arguments are standard for proving correctness of compilers for deterministic languages. However, the situation here is more complex due to the fact that the target CFG has a much wider space of traces: the values of each versioned variable in the target program are initially unconstrained, meaning traces exist for all of their combinations. On the other hand, many of these traces do not survive the **assume** statements encountered in the target program. Picking the correct *single* trace or state to simulate a particular source execution would require knowledge of all variable assignments that are *going* to happen, which

is not possible due to non-determinism and would preclude the block-modular proof strategies that our validation approach employs.

Instead, we generalise this idea to relating each single source state s with a set T of corresponding target program states. We define variable relations \mathcal{V}_R at each point in a trace, making explicit the mappings used in the SSA step between source program variables and their corresponding versions. For example, on entry to block B'_2 in the source version of our running example (correspondingly B''_2 in the target), the \mathcal{V}_R relation relates i to $i1$ and j to $j2$. All states $t \in T$ must precisely agree with s w.r.t. \mathcal{V}_R (e.g., $s(i) = t(i1)$, $s(j) = t(j2)$). On the other hand, our sets of states T are defined to be completely unconstrained (besides typing) for *future* variable versions. For example, for every $t \in T$ at the same point in our example, there will be states in T assigning each possible value (of the same type) to $i2$ (and otherwise agreeing with t).

More precisely, for a set of variables X , we say that a set of states T *constrains at most* X w.r.t. variable context Λ if, for every $t \in T$, $z \notin X$, z is in Λ , and value v of z 's type, we have $t[z \mapsto v] \in T$. In other words, the set T is closed under arbitrary changes to values of all variables in Λ but *not* in X . We construct our sets T such that they constrain at most *current and past versions* of program variables. It is this fact that enables us to handle subsequent **assume** statements in the target program and, in particular, to show that the set of possible traces in the target program never becomes empty while there are possible traces in the source program. For example, when relating the source command $j := j+1$ in B'_3 with the target command **assume** $j3 = j2 + 1$ in block B''_3 , we use the fact that our set of states does not constrain $j3$ to prove that, although many traces go to magic at this point, for a non-empty set of states $T' \subseteq T$ (those in which $j3$ has the “right” value equal to $j2 + 1$), execution continues in the target.

We now make these notions more precise by showing the definition of our local block lemmas for the passification phase⁹.

Theorem 2 (Passification Local Block Lemma). *Let B be a source block with commands cs , whose corresponding target block has commands cs' ; let \mathcal{V}_R and \mathcal{V}'_R be the variable relations at the beginning and end of B , respectively. Let X be a set of variable versions, and $N(ns)$ be a normal state. Let T be a non-empty set of normal states such that $N(ns)$ agrees with T according to \mathcal{V}_R , and T constrains at most X w.r.t. Λ_2 . Furthermore, let Y be the variable versions corresponding to the targets of assignment and havoc statements in cs . If both*

1. $A, \Lambda_1, \Gamma, \Omega \vdash \langle cs, N(ns) \rangle [\rightarrow] s' \wedge s' \neq M$
2. $X \cap Y = \emptyset$

then there exists a non-empty set of normal states $T' \subseteq T$ s.t. T' constrains at most $X \uplus Y$ w.r.t. Λ_2 and for each $t' \in T'$, there exists a state t'^ s.t.*

1. $A, \Lambda_2, \Gamma, \Omega \vdash \langle cs_2, t' \rangle [\rightarrow] t'^* \wedge (s' = F \implies t'^* = F)$
2. *If s' is a normal state, then s' and t' are related w.r.t. \mathcal{V}'_R (and $t'^* = t'$).*

⁹ We omit some details regarding well-typedness, handled fully in our formalisation.

This lemma captures our generalised notion of forward simulation appropriately. The first conclusion expresses that the target does not get stuck and that failures are preserved, while the second shows that if the source execution neither fails nor stops then the resulting states are related. Note that premise 2 is essential in the proof to guarantee that the **assume** statements introduced by passification do not eliminate the chance to simulate source executions; the condition expresses that the variable versions newly constrained do not intersect with those previously constrained. To prove these lemmas over the commands in a single block, we are forced to check that the same version is not constrained twice.

5.3 Passification Certification: Global Block Theorems

As for all phases, we lift our local block lemmas to theorems certifying all executions *starting* from a particular block, and thus, ultimately, to entire CFGs. For the passification phase, most of the conceptual challenges are analogous to those of the local block lemmas; we similarly employ \mathcal{V}_R relations between source variables and their corresponding target versions. To connect with our local block lemmas (and build up our global block theorems, which we do backwards through the CFG structure), we repeatedly require the key property that the set of variable versions constrained in our executions so far is disjoint from those which may be constrained by a subsequent **assume** statement (*cf.* premise 2 of our local block lemma above). Concretely tracking and checking disjointness of these concrete sets of variables is simple, but turns out to get expensive in Isabelle when the sets are large.

We circumvent this issue with our own *global versioning scheme* (as opposed to the versions used by Boogie, which are *independent* for different source variables): according to the CFG structure, we assign a *global* version number $\text{ver}_G(x)$ to each variable x in the target program such that, if x is constrained in a target block B' and y is constrained in another target block B'' reachable from B' , then $\text{ver}_G(x) < \text{ver}_G(y)$. Such a consistent global versioning always exists in the target programs generated by Boogie because the only variables not constrained exactly once *in the program* are those used to synchronise executions (i.e. `j4` in Fig. 5), which always appear right before branches are merged. We can now encode our disjointness properties much more cheaply: we simply compare the *maximal* global version of all already-constrained variables with the *minimal* global version of those (potentially) to be constrained. Since we represent variables as integers in the mechanisation, we directly use our global version *as* the variable name for the target program; there is no need for an extra lookup table. Note that (readability aside) it makes no difference which variables names are used in intermediate CFGs; we ultimately care only about validating the original CFG.

6 The VC Phase

In this section, we present the validation of the VC phase in the Boogie verifier. This phase has two main aspects: (1) it encodes and desugars all aspects of the

Boogie type system, employing additional uninterpreted functions and axioms to express its properties [33]; program expression elements such as Boogie functions are analogously desugared in terms of these additional uninterpreted functions, creating a non-trivial logical gap between expressions as represented in the VC and those from the input program. (2) It performs an efficient (block-by-block) calculation of a weakest precondition for the (acyclic, passified) CFG, resulting in a formula characterising its verification requirements, subject to background axioms and other hypotheses.

6.1 VC Structure

The generated VC has the following overall structure (represented as a shallow embedding in our certificates)¹⁰:

$$\forall \underbrace{VC \text{ quantifiers}}_{\text{type encoding parameters, functions, variable values}} . \left(\underbrace{VC \text{ assumptions}}_{\text{type encoding, func./var./prog. axioms}} \implies CFG \text{ WP} \right)$$

The VC quantifies over parameters required for the type encoding, as well as VC counterparts representing the variable values and functions in the Boogie program. The VC body is an implication, whose premise contains: (1) assumptions that axiomatise the type encoding parameters, (2) axioms expressing the typing of Boogie variables and functions, and (3) assumptions directly relating to axioms explicitly declared in the Boogie program. The conclusion of the implication is an optimised version of the weakest (liberal) precondition (WP) of the CFG.¹¹

6.2 Boogie’s Logical Encoding of the Boogie Type System

We first briefly explain Boogie’s logical encoding of its own type system. Values and types are represented at the VC level by two uninterpreted carrier sorts V and T . An uninterpreted function typ from V to T maps each value to the representation of its type. Boogie type constructors are each modelled with an (injective) uninterpreted function C with return sort T and taking arguments (per constructor parameter) of sort T . For example, a type constructor $List(t)$ is represented by a VC function from T to T . Projection functions are also generated for each type constructor (C_i^T for each type argument at position i), e.g. mapping the representation of a type $List(t)$ to the representation of type t .

This encoding is then used in the VC to recover Boogie typing constraints for the untyped VC terms. Recovering the constraints is not always straightforward due to optimisations performed by Boogie. For example, the VC translation of the Boogie expression $\forall_{ty} t. \forall x : List(t). e$ no longer quantifies over types;

¹⁰ Note that top-level quantification over functions is implicit in the (first-order) SMT problem generated by Boogie; we quantify explicitly in our Isabelle representation.

¹¹ One difference in our version of the Boogie verifier is that we switched off the generation of extra variables introduced to report error traces [32]; these are redundant for programs that do not fail and further complicate the VC structure.

all original occurrences of t in e having been translated to $List_1^\pi(typ(x))$. This optimisation reflects that this particular type quantification is redundant, since t can be recovered from the type of x .¹²

6.3 Working from VC Validity

Our certificates assume that the generated VC is valid (certifying the validity-checking of the VC by an SMT solver is an orthogonal concern). However, connecting VC validity back to block-level properties about the specific program requires a number of technical steps. We need to construct Isabelle-level semantic values to *instantiate* the top-level quantifiers in the VC such that the corresponding VC assumptions (left-hand side of the VC) can be proved and, thus, validity of the corresponding WP can be deduced. Moreover, we must ensure that our instantiation yields a WP whose validity implies correctness of the Boogie program. For example, a top-level VC quantifier modelling a Boogie function f must be instantiated with a mathematical function that behaves in the same way as f for arguments of the correct type.

We instantiate the carrier sort V for values in the VC with the corresponding type denoting Boogie values in our formalisation; the carrier sort T for *types* is instantiated to be all Boogie types that do not contain free variables (i.e. closed types). Constructing explicit models for the quantified functions used to model Boogie’s type system (satisfying, e.g., suitable inverse properties for the projection functions) is straightforward. For the VC-level variable values, we can directly instantiate the corresponding values in the initial Boogie program state.

VC-level functions representing those declared in the Boogie program are instantiated as (total) functions which, *for input values of appropriate type* (the arguments and output are untyped values of sort V), are defined simply to return the same values as the corresponding function in our model. However, perhaps surprisingly, Boogie’s VC embedding of functions logically requires functions to return values of the specified return type even if the input values do not have the types specified by the function. In such cases, we define the instantiated function to return some value of the specified type, which is possible since in well-formed contexts every closed type has at least one value in our model.

After our instantiation, we need to prove the hypotheses of the VC’s implication; in particular that all axioms (both those generated by the type system encoding and those coming from the program itself) are satisfied. The former are standard and simple to prove (given the work above), while the latter largely follow from the assumption that each declared axiom must be satisfied in the initial state restricted to the constants. The only remaining challenge is to relate VC expressions with the evaluation of corresponding Boogie expressions; an issue which also arises (and is explained) below, where we show how to connect validity of the instantiated WP to the program.

¹² Note that in the VC the quantification over x ranges over all values of sort V . An implication is used to consider only those x for which $typ(x) = List(List_1^\pi(typ(x)))$.

6.4 Certifying the VC Phase

Boogie’s weakest precondition calculation is made size-efficient by the usage of explicit named constants for the weakest preconditions $wp(B, \mathbf{true})$ for each block B , which is defined in terms of the named constants for its successor blocks. For example, in Fig. 5, $wp(B_2'', \mathbf{true})$ is given by $i_1^{vc} \neq 0 \implies wp(B_3'', \mathbf{true}) \wedge wp(B_4'', \mathbf{true})$. Here i_1^{vc} is the value that we instantiated for the variable `i1`.

We exploit this modular construction of the generated weakest precondition for the local and global block theorems. We prove for each block B with commands cs the following local block lemma:

Theorem 3 (VC Phase Local Block Lemma).

If $A, \Lambda, \Gamma, \Omega \vdash \langle cs, N(ns) \rangle [\rightarrow] s'$ and $wp(B, \mathbf{true})$ holds, then $s' \neq F$ and if s' is a normal state, then $\forall B_{suc} \in \text{successors}(B). wp(B_{suc}, \mathbf{true})$.

Once one has proved this lemma for all blocks in the CFG, combining them to obtain the corresponding global block theorems (via our usual reverse walk of the CFG) is straightforward. The main challenge is in decomposing the proof for the local block lemma itself for a block B , for which we outline our approach next.

By this phase, the first command in B must be either an **assume** e or an **assert** e command. In the former case, we rewrite $wp(B, \mathbf{true})$ into the form $e^{vc} \implies H$, where e^{vc} is the VC counterpart of e and where H corresponds to the weakest precondition of the remaining commands. This rewriting may involve undoing certain optimisations Boogie’s implementation performed on the formula structure. Next, we need to prove that e evaluates to e^{vc} (see below). Hence, if e evaluates to **true** (the execution does not go to magic) then H must be true, and we can continue inductively. The argument for **assert** e is similar but where we rewrite the VC to $e^{vc} \wedge H$ (i.e. e^{vc} and H must both hold); if e evaluates to e^{vc} , we know that the execution does not fail.

Proving that e evaluates to e^{vc} arises in both cases and also in our previous discharging of VC hypotheses. Note that, in contrast to e , e^{vc} is not a Boogie expression, but a shallowly embedded formula that includes the instantiations of quantified variables we constructed above. Showing this property works largely on syntax-driven rules that relate a Boogie expression with its VC counterpart, except for extra work due to mismatching function signatures and optimisations that Boogie made either to the formula structure or via the type system encoding (cf. Sec. 6.2). We handle some of these cases by showing that we can rewrite the formula back into the unoptimised standard form we require for our syntax-driven rules and in other cases we directly work with the optimised form. Both cases are automated using Isabelle tactics.

This concludes our discussion of the certification of Boogie’s three key phases. Combining the three certificates yields an end-to-end proof that the validity of the generated verification conditions implies the correctness of the input program, that is, that the given verification run is sound.

Table 1. Selection of algorithmic examples with the lines of code (LOC), the number of procedures (#P), the time it takes for Isabelle to check the certificate in seconds (the average of 5 runs on a Lenovo T480 with 32 GB, i7-8550U 1.8 GhZ, Ubuntu 18.04 on the Windows Subsystem for Linux), and the certificate size expressed as the number of non-empty lines of Isabelle.

Name	LOC	#P	Time [s]	Size
TuringFactorial	29	1	19.4	1986
Find	27	2	27.3	2100
DivMod	69	2	28.4	4753
Summax [27]	23	1	19.1	1953
MaxOfArray [12]	22	1	19.9	1944
SumOfArray [12]	22	1	18.7	1534
Plateau [12]	50	1	22.9	2019
WelfareCrook [12]	52	1	39.4	2528
ArrayPartitioning [12]	57	2	27.6	3514
DutchFlag [12]	76	2	52.8	3994

7 Implementation and Evaluation

In this section, we evaluate our certifying version of the Boogie verifier [36], which produces Isabelle certificates proving the correctness of Boogie’s pipeline for programs it verifies.

We have implemented our validation tool as a new C# module compiled with Boogie. We instrumented Boogie’s codebase to call out to our module, which allows us to obtain information that we can use to validate the key phases, and extended parts of the codebase to extract information more easily. Moreover, we disabled counter-example related VC features and the generation of VC axioms for any built-in types and operators that we do not support. We added or changed fewer than 250 non-empty, uncommented lines of code across 11 files in the existing Boogie implementation.

Given an input file verified by Boogie, our work produces an Isabelle certificate per procedure p that certifies the correctness of the corresponding CFG-to-DAG source CFG as represented internally in Boogie. The generation and checking of the certificate is fully automatic, without any user input. We use a combination of custom and built-in Isabelle tactics. In addition to the three key phases we describe in detail, our implementation also handles several smaller transformations made by Boogie, such as constant propagation. Our tool currently supports the default options of Boogie (only) and does not support advanced source-level *attributes* (for instance, to selectively force procedures to be inlined).

We evaluated our work in two ways. Firstly, to evaluate the applicability of our certificate generation, we automatically collected all input files with at least one procedure from Boogie’s test suite [1] which verify successfully and which either use no unsupported features or are easily desugared (by hand) into versions without them. This includes programs with procedure calls since Boogie simply desugars these in an early stage. For programs employing attributes, we

checked whether the program still verifies *without* attributes, and if so we also kept these. In total, this yields 100 programs from Boogie’s test suite. Secondly, we collected a corpus of ten Boogie programs which verify interesting algorithms with non-trivial specifications: three from Boogie’s test suite and seven from the literature [12, 27]. Where needed we manually desugared usages of Boogie maps (which we do not yet support) using type declarations, functions, and axioms.

Of the 100 programs from Boogie’s test suite, we successfully generate certificates in 96 cases. The remaining 4 cases involve special cases that we do not handle yet. For 2 of them, extending our work is straightforward: one special case includes a naming clash and the other case can be amended by using a more specific version of a helper lemma. The remaining two fail because of our incomplete handling of function calls in the VC phase when combined with coercions between VC integers or booleans and their Boogie counterparts. Handling this is more challenging but is not a fundamental issue.

For the corpus of 10 examples, Tab. 1 shows the generated certificate size and the time for Isabelle to check their validity.¹³ The ratio of certificate size to code size ranges from 41 to 89; this rather large ratio emphasises the substantial work in formally validating the substantial work which Boogie’s implementation performs. Optimisations to further reduce the ratio are possible. The validation of certificates takes usually under one second per line of code. While these times are not short, they are acceptable since certificate generation needs to run only for (verified) release versions of the program in question.

8 Related Work

Several works explore the validation of program verifiers. Garchery et al. [20] validate VC rewritings in the Why3 VC generator [16]. Unlike our work, they do not connect VCs with programs and do not handle the erasure of polymorphic types. Strub et al. [39] validate part of a previous version of the F* verifier [40] by generating a certificate for the F* type checker itself, which type checks programs by generating VCs. Like us, they assume the validity of the generated VC itself, but they do not consider program-to-program transformations such as ours. Another approach is taken by Aguirre [2] who shows how one can map proofs of the VC back to correctness of an F* program. They prove a once-and-for-all result, but the approach could be lifted to a validation approach using the proof-producing capability of SMT solvers [7]. Lifting the approach would require extending the work to handle classical instead of constructive VC proofs.

There is some work on proving VC generator implementations correct once and for all, although none of the proven tools are used in practice. Homeier and Martin [23] prove a VC generator correct in HOL for an executable language and a simpler VC phase than Boogie’s. Herms et al. [22] prove a VC generator inspired by Why3 correct in Coq. However, some more-challenging aspects of Why3’s VC transformation and polymorphic type system are not handled. Vogels

¹³ The time to generate the certificate is not included, but is negligible here.

et al. [44] prove a toolchain for a Boogie-like language correct in Coq, including passification and VC phases. However, the language is quite limited: without unstructured control flow, loops (i.e. no need for a CFG-to-DAG phase), functions, or polymorphism (i.e. no type encoding). Verifiers other than VC generators, include the verified Verasco static analyzer [25], which supports a realistic subset of C, but whose performance is not yet on par with unverified, industrial analyzers.

Validation has also been explored in other settings. Alkassar et al. [3] adjust graph algorithms to produce witnesses that can be then used by verified validators to check whether the result is correct. In the context of compiler correctness, many validation techniques express a per-run validator in Coq, prove it correct once-and-for-all [8, 41, 43], and then extract executable code (the extraction must be trusted). In the verified CompCert compiler [34], such validators have been used in combination with the once-and-for-all approach. Validators are used for phases that can be more easily validated than proved correct once and for all. One such example related to our certification of the passification phase is the validation of the SSA phase [8], dealing also with versioned variables in the target (but not with **assume** statements that prune executions). In contrast to our work, they require an explicit notion of CFG domination and they do not use a global versioning scheme to efficiently check that two parts of the CFG constrain disjoint versions. Our versioning idea is similar to a technique used for the validation of a dominator relation in a CFG [9], which assigns intervals to basic blocks (as opposed to assigning versions to variables) to efficiently determine whether a block dominates another one. The validation of the Cogent compiler [38] follows a similar approach to ours in that it generates proofs in Isabelle.

9 Conclusion

We have presented a novel verifier validation approach, and applied it successfully to three key phases of the Boogie verifier, providing formal underpinnings for both the language and its verifier for the first time. Our work demonstrates that it is feasible to provide strong formal guarantees regarding the verification results of practical VC generators written in modern mainstream languages.

In the future, we plan to extend our supported subset of Boogie, e.g. to include procedure calls and bitvectors. Supporting Boogie’s potentially-impredicative maps is the main open challenge: maps can take other maps as input, potentially including themselves. The challenge with this feature is to still be able to express a type in Isabelle capturing all Boogie values despite the potentially-cyclic nature of map types. In practice, however, this may not be required in full generality: we have observed that Boogie front-ends rarely use maps that contain maps of the same type as input. Therefore, we plan to extend our technique to support a suitably-expressive restricted form of Boogie maps.

Acknowledgements. We thank Alain Delaët–Tixeuil for his earlier work on this topic, Thibault Dardinier for improving our artifact, Martin Clochard for helpful discussions and the anonymous reviewers for their valuable comments.

This work was partially funded by the Swiss National Science Foundation (SNSF) under Grant No. 197065.

References

1. Boogie verifier repository. <https://github.com/boogie-org/boogie>
2. Aguirre, A.: Towards a provably correct encoding from F* to SMT. Tech. rep., INRIA (2016)
3. Alkassar, E., Böhme, S., Mehlhorn, K., Rizkallah, C.: A framework for the verification of certifying computations. *JAR* **52**(3), 241–273 (2014)
4. Astrauskas, V., Müller, P., Poli, F., Summers, A.J.: Leveraging Rust types for modular specification and verification. In: *OOPSLA* (2019)
5. Barnett, M., Fähndrich, M., Leino, K.R.M., Müller, P., Schulte, W., Venter, H.: Specification and verification: The Spec# experience. *CACM* **54**(6), 81–91 (June 2011)
6. Barnett, M., Leino, K.R.M.: Weakest-precondition of unstructured programs. In: *PASTE* (2005)
7. Barrett, C., de Moura, L., Fontaine, P.: Proofs in satisfiability modulo theories. In: *All about Proofs, Proofs for All, Mathematical Logic and Foundations*, vol. 55, pp. 23–44. College Publications (2015)
8. Barthe, G., Demange, D., Pichardie, D.: Formal verification of an SSA-based middle-end for compcert. *TOPLAS* **36**(1) (2014)
9. Blazy, S., Demange, D., Pichardie, D.: Validating dominator trees for a fast, verified dominance test. In: *ITP* (2015)
10. Blom, S., Darabi, S., Huisman, M., Oortwijn, W.: The VerCors tool set: Verification of parallel and concurrent software. In: *iFM* (2007)
11. Böhme, S., Weber, T.: Fast LCF-style proof reconstruction for Z3. In: *ITP* (2010)
12. Chen, Y., Furia, C.A.: Triggerless happy – intermediate verification with a first-order prover. In: *iFM* (2017)
13. Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A practical system for verifying concurrent C. In: *TPHOLs* (2009)
14. Coq Development Team, T.: The Coq Reference Manual, version 8.10 (2019), available electronically at <http://coq.inria.fr/documentation>
15. Ekici, B., Mebsout, A., Tinelli, C., Keller, C., Katz, G., Reynolds, A., Barrett, C.W.: SMTCoq: A plug-in for integrating SMT solvers into Coq. In: *CAV* (2017)
16. Filliâtre, J.C., Paskevich, A.: Why3 — where programs meet provers. In: *ESOP* (2013)
17. Filliâtre, J.C., Marché, C.: The Why/Krakatoa/Caduceus platform for deductive program verification. In: *CAV* (2007)
18. Flanagan, C., Saxe, J.B.: Avoiding exponential explosion: Generating compact verification conditions. In: *POPL* (2001)
19. Fleury, M., Schurr, H.: Reconstructing veriT proofs in Isabelle/HOL. In: *PxTP* (2019)
20. Garchery, Q., Keller, C., Marché, C., Paskevich, A.: Des transformations logiques passent leur certificat. In: *JFLA* (2020)
21. Hecht, M.S., Ullman, J.D.: Flow graph reducibility. *SIAM J. Comput.* **1**(2), 188–202 (1972)

22. Herms, P., Marché, C., Monate, B.: A certified multi-prover verification condition generator. In: VSTTE (2012)
23. Homeier, P.V., Martin, D.F.: A mechanically verified verification condition generator. *The Computer Journal* **38**(2), 131–141 (1995)
24. Isabelle Development Team, T.: The Isabelle Documentation, version June 2019 (2019), available electronically at <https://isabelle.in.tum.de/documentation.html>
25. Jourdan, J.H., Laporte, V., Blazy, S., Leroy, X., Pichardie, D.: A formally-verified C static analyzer. In: POPL (2015)
26. Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: Frama-c: A software analysis perspective. *Formal Aspects of Computing* **27**(3), 573–609 (2015)
27. Klebanov, V., Müller, P., Shankar, N., Leavens, G.T., Wüstholtz, V., Alkassar, E., Arthan, R., Bronish, D., Chapman, R., Cohen, E., Hillebrand, M., Jacobs, B., Leino, K.R.M., Monahan, R., Piessens, F., Polikarpova, N., Ridge, T., Smans, J., Tobies, S., Tuerk, T., Ulbrich, M., Weiß, B.: The 1st verified software competition: Experience report. In: FM (2011)
28. Lal, A., Qadeer, S., Lahiri, S.K.: A solver for reachability modulo theories. In: CAV (2012)
29. Leino, K.R.M.: This is Boogie 2 (June 2008), <https://www.microsoft.com/en-us/research/publication/this-is-boogie-2-2/>
30. Leino, K.R.M.: Efficient weakest preconditions. *Inf. Process. Lett.* **93**(6), 281–288 (2005)
31. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: LPAR (2010)
32. Leino, K.R.M., Millstein, T.D., Saxe, J.B.: Generating error traces from verification-condition counterexamples. *Science of Computer Programming* **55**(1-3), 209–226 (2005)
33. Leino, K.R.M., Rümmer, P.: A polymorphic intermediate verification language: Design and logical encoding. In: TACAS (2010)
34. Leroy, X.: Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In: POPL (2006)
35. Müller, P., Schwerhoff, M., Summers, A.J.: Viper: A verification infrastructure for permission-based reasoning. In: VMCAI (2016)
36. Parthasarathy, G., Müller, P., Summers, A.J.: Formally validating a practical verification condition generator – artifact (2021). <https://doi.org/10.5281/zenodo.4726554>
37. Parthasarathy, G., Müller, P., Summers, A.J.: Formally validating a practical verification condition generator (extended version). arXiv:2105.14381 (2021)
38. Rizkallah, C., Lim, J., Nagashima, Y., Sewell, T., Chen, Z., O’Connor, L., Murray, T., Keller, G., Klein, G.: A framework for the automatic formal verification of refinement from Cogent to C. In: ITP (2016)
39. Strub, P.Y., Swamy, N., Fournet, C., Chen, J.: Self-certification: Bootstrapping certified typecheckers in F* with Coq. In: POPL (2012)
40. Swamy, N., Hritcu, C., Keller, C., Rastogi, A., Delignat-Lavaud, A., Forest, S., Bhargavan, K., Fournet, C., Strub, P.Y., Kohlweiss, M., Zinzindohoué, J.K., Zanella-Béguelin, S.: Dependent types and multi-monadic effects in F*. In: POPL (2016)
41. Tristan, J.B., Leroy, X.: Formal verification of translation validators: A case study on instruction scheduling optimizations. In: POPL (2008)
42. Tristan, J.B., Leroy, X.: Verified validation of lazy code motion. In: PLDI (2009)
43. Tristan, J.B., Leroy, X.: A simple, verified validator for software pipelining. In: POPL (2010)
44. Vogels, F., Jacobs, B., Piessens, F.: A machine-checked soundness proof for an efficient verification condition generator. In: SAC (2010)