

Building a Push-Button RESOLVE Verifier: Progress and Challenges

Murali Sitaraman, Bruce Adcock, Jeremy Avigad, Derek Bronish, Paolo Bucci, David Frazier, Harvey M. Friedman, Heather Harton, Wayne Heym, Jason Kirschenbaum, Joan Krone, Hampton Smith, and Bruce W. Weide

Technical Report RSRG-09-01

School of Computing
100 McAdams
Clemson University
Clemson, SC 29634-0974 USA

January 2009

Copyright © 2009 by the authors. All rights reserved.

Building a Push-Button RESOLVE Verifier: Progress and Challenges

Murali Sitaraman¹, Bruce Adcock², Jeremy Avigad³, Derek Bronish², Paolo Bucci², David Frazier¹, Harvey M. Friedman⁴, Heather Harton¹, Wayne Heym², Jason Kirschenbaum², Joan Krone⁵, Hampton Smith¹, and Bruce W. Weide²

¹ School of Computing, Clemson University, Clemson, SC 29634, USA
{murali, hamptos, hkeown, dfrazie}@cs.clemson.edu

² Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA
{adcockb, bucci, bronish, heym, kirschen, weide}@cse.ohio-state.edu

³ Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA 15213 USA
avigad@cmu.edu

⁴ Department of Mathematics, The Ohio State University, Columbus, OH 43210, USA
friedman@math.ohio-state.edu

⁵ Department of Mathematics and Computer Science, Denison University, Granville, OH 43023, USA
krone@denison.edu

Abstract. A central objective of the verifying compiler grand challenge is to develop a push-button verifier that generates proofs of correctness in a syntax-driven fashion similar to the way an ordinary compiler generates machine code. The software developer's role is then to provide suitable specifications and annotated code, but otherwise to have no direct involvement in the verification step. The general mathematical developments and results upon which software correctness is based are established through a separate formal proof process. Current versions of the RESOLVE verifier generate verification conditions (VCs) for the correctness of component-based software in a modular fashion — one component at a time. The VCs are currently verified using automated capabilities of the Isabelle proof assistant and some specialized decision procedures. Initial experiments with the tools and further analytic considerations show both the progress that has been made and the challenges that remain.

1 Introduction

Automation is a fundamental challenge in developing a verifier and eventually realizing the grand challenge of building verified software [25][26]. In pursuit of automation, some well-known efforts have focused on verifying particular sets of properties of software. Alternatively, to achieve full behavioral verification, others have restricted the domain of verification to types that are typically “built in” to programming languages, such as integers and arrays. In the interest of achieving automation, still others have proposed systems that, by design, compromise on soundness and/or completeness.

The objective of our verifying compiler work is to facilitate specification and push-button verification of full behavior of software. It is designed to verify software written in RESOLVE, an integrated specification and programming language with clean semantics that has been especially designed to facilitate verification. The specification language is a mathematical higher-order language and is extensible with new mathematical theories, definitions, and related developments including general-purpose theorems and lemmas that are proved off-line once and for all. The implementation language is imperative and object-based. The proof system for verification of implementation correctness is both sound and relatively complete. The proof system is also modular: it is possible to verify component-based software, one component at a time, relying on only the specifications of the components reused by the component that is being verified. The goal of push-button verification is that beyond providing specifications and implementations with suitable annotations (e.g., loop invariants and progress metrics), software developers should not be involved in proof activities; proofs of non-trivial theorems used in program verification, however, might themselves require assistance from mathematical specialists, but that is outside the push-button process.

This paper explains the progress we have made in developing a verifier and related tools, and outlines several challenges that remain. Its focus is on tools not on the underlying theory (for which relevant background papers are cited); the paper is intended to be somewhat self-contained and hence accessible to a reader who is not familiar with these details. Section 2 of the paper presents the objectives of the verifying compiler through illustrative screenshots of a tool demonstration system for a simple benchmark. It also describes an internal architecture of the verifier that serves to integrate the subordinate tools detailed in subsequent sections. Section 3 explains VC generation using two versions that we have developed to explore issues in generation of VCs and their proofs; for variety, the section employs two distinct examples, both of which illustrate the modular nature of VC

generation. The focus of Section 4 is on proving these VCs automatically in two different ways. One employs Isabelle [41] as an automated prover (as opposed to an interactive proof assistant) and another employs a simplifier and specialized decision procedure we have developed. Section 5 summarizes tool-related issues in checking proofs of non-trivial theorems invoked in completing automated proofs as in Section 4. Section 6 contains a discussion of related automated software verification efforts along with our conclusions.

2 Verifying Compiler Overview

This section presents an overview of the verifying compiler and a description of the toolset underlying its overall internal organization.

2.1 External Vision of Push-Button Verification

To quickly illustrate the goals of the verifier, we consider the first automated verification benchmark described in [48]. This benchmark requires verification of (total) correctness of an operation to add two integers by repeated incrementing and decrementing. The implementation may be annotated as necessary with invariants for loops and/or progress metrics to show termination.

Three screenshots of our tool demonstration system illustrate the working of one version of the RESOLVE tools for this simple example, explained in detail in [48]. Figure 1 contains the specification of *Integer* addition within a larger context of the *Integer_Template* concept (or contract) that is not shown. Specification of the *Integer_Template* concept, or any other one selected, may be viewed by clicking the “View” button. The picture also lists two implementations of *Integer* addition. Figure 2 shows an iterative implementation of the specification. Figure 3 shows the result from the push-button verifier on clicking the “Generate and Verify” button. (The “Compile” button generates executable code and its role is obvious.)

The example illustrates that given supporting specifications and relevant mathematics, the task of the programmer is limited to supplying annotated code for the implementation under consideration for verification. All detailed proof activity is left to the verifier.

Verification Demo		Math Units	
Concept	Enhancement	Concept Realization	Enhancement Realization
Location_Linking_Template.co List_Template.co Integer_Template.co Char_Str_Template.co Basic_Map_Template.co Alt_Queue_Template.co	Int_Mult_Capability.en Adding_Capability.en		Recursive_Add_to_Realiz.rb Iterative_Add_to_Realiz.rb
View	View	View Compile	View

Code for Adding_Capability.en enhancement:

```

Enhancement Adding_Capability for Integer_Template:
  Operation Add_to(updates i:Integer; evaluates j:Integer);
  requires min_int <= i + j and i + j <= max_int and j >= 0;
  ensures i = #i + j;
end Adding_Capability;
  
```

Generate Code

Generate Isabelle VCs

Generate Readable VCs

Generate and Verify

Figure 1: Example Specification: Integer Addition

Verification Demo		Math Units	
Concept	Enhancement	Concept Realization	Enhancement Realization
Prioritizer_Template.co Location_Linking_Template.co List_Template.co Integer_Template.co Char_Str_Template.co Basic_Map_Template.co	Int_Mult_Capability.en Adding_Capability.en		Recursive_Add_to_Realiz.rb Iterative_Add_to_Realiz.rb
View	View	View Compile	View

Code for Iterative_Add_to_Realiz.rb enhancement realization:

```

Realization Iterative_Add_to_Realiz for Adding_Capability of Integer_Template:
  uses Std_Boolean_Fac;

  Procedure Add_to(updates i:Integer; evaluates j:Integer);
  While (not Is_Zero(j))
    changing i, j;
    maintaining (i + j = #i + #j) and j >= 0;
    do
      decreasing j;
      Increment (i);
      Decrement (j);
    end;
  end Add_to;
end Iterative_Add_to_Realiz;
  
```

Generate Code

Generate Isabelle VCs

Generate Readable VCs

Generate and Verify

Figure 2: An Iterative Implementation of Addition

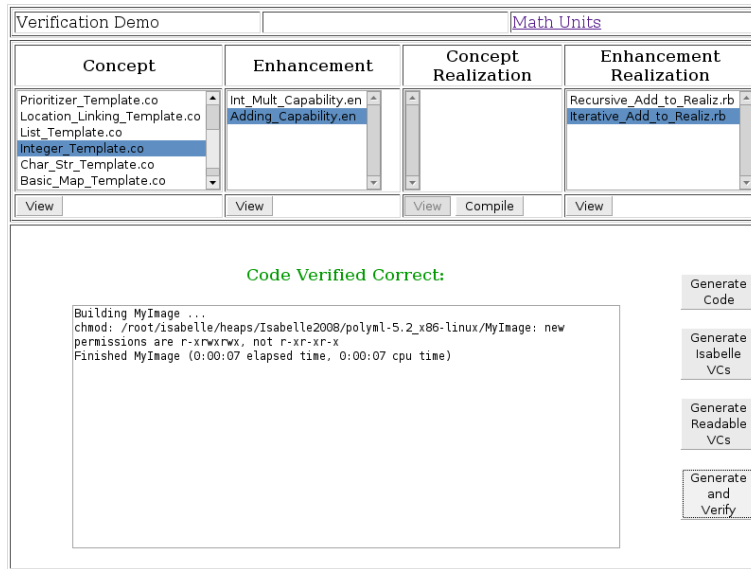


Figure 3: Net Result of Push-button Verification

The verification process is *modular*, so verification of one implementation never involves any other implementations; merely other specifications. In other words, every implementation is verified correct only once in its lifetime. For example, to complete benchmark #1, part 2[48] where the goal is to verify Integer multiplication implemented using repeated addition, only the specification of *Integer* addition operation (in Figure 1) is necessary. The implementations of addition are not relevant and are not used.

2.2 Toolset Organization and Its Usage Scenarios

Figure 4 is an illustration of the internal organization of the verifier. In the figure, we depict tangible artifacts in the verified software process by ovals, and tools that facilitate the process by boxes. Two human roles are shown as shaded boxes; it should be noted that both roles might be filled by the same person, or each by teams of people. The first role is to write “assertive code” in accordance with specifications of desired software behavior. The second role is to develop mathematical theories, definitions and results relevant to the specifications, along with mathematical proofs of such purely mathematical results, proofs that may be

generated automatically or written manually, but validated by a mechanical proof checker. But proofs of correctness of the code are to be completed automatically with no human intervention.

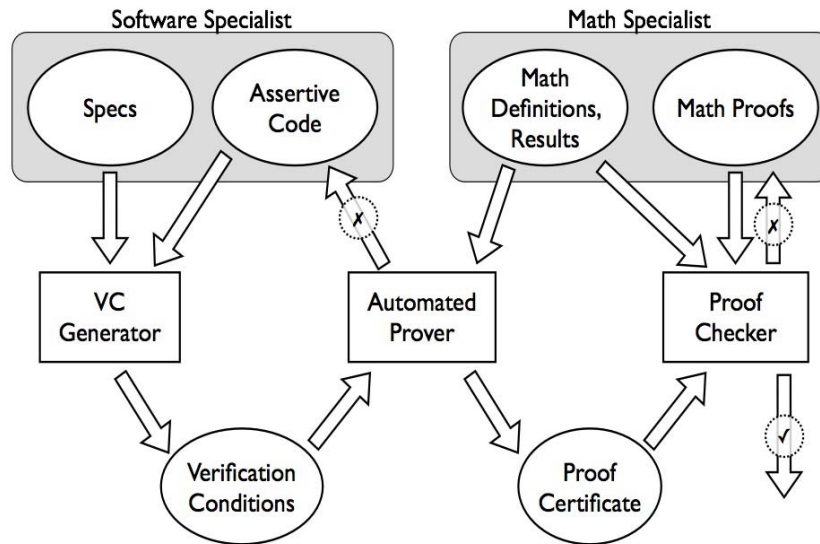


Figure 4: Verification Toolset Architecture

A person in the first role—we call her the *software specialist*—writes not only the kind of (executable) code routinely written today, but also embedded formal annotations such as invariants and progress metrics for loops, representation invariants and abstraction relations for new data representations, etc. She debugs and repairs code and/or assertions when an implementation cannot be automatically verified. In addition to understanding specifications of off-the-shelf components used in her code, the software specialist also writes specifications for local operations she needs in client code, and designs, specifies, and implements more general-purpose components for shared component libraries.

When there is difficulty formalizing intended software behavior, a software specialist may consult a person in the second role—the *math specialist*. The math specialist may engage in new theory development or add mathematical definitions and results to existing theories to simplify specifications and code annotations, and may establish various lemmas and theorems about fragments of mathematics to ease the task of the automated prover. For example, when writing code involving

text, stacks, queues, lists, and related components (which in RESOLVE are modeled using mathematical strings), it is sometimes necessary to reason about the reversal of a string as can be seen in Section 3.1. The math specialist might define a mathematical function *reverse* from strings to strings, and then write lemmas that might be expected to arise in the verification of programs involving types that are modeled as strings. As an example, consider the lemma $\forall \alpha, \beta (reverse(\alpha \bullet \beta) = reverse(\beta) \bullet reverse(\alpha))$, where \bullet denotes concatenation of strings. A key idea in the verified software paradigm is that the proof of such a supporting mathematical fact need not be carried out automatically. The math specialist does this kind of proof off-line, perhaps with the help of an automated proof assistant such as Isabelle. Because the lemma being proved is assumed throughout the rest of the software verification process, the soundness of the system depends on the validity of this proof; so submitting the proof to an automated proof checker may be appropriate as explained in Section 5. The proved lemma then becomes one of many mathematical results included with the definitions used in specifications and becomes available to the automated prover during software verification, where the proof of the lemma is no longer involved.

Some lemmas might require substantial mathematical insight to identify and prove, so the goal for push-button verification of most programs *given* a sufficiently rich set of theories and developments in the form of lemmas in the automated prover's math library. In the steady state, it is expected that the role of the math specialist becomes not primarily to add more and more lemmas for established theories, but rather to develop additional theories and results that might be helpful in new software development tasks for new application domains.

It is likely the verified software paradigm will involve far more software specialists than math specialists, if only because the work of the latter is inherently far more reusable. Considerable experience to date with this paradigm suggests there will be far fewer definitions and results in the library of mathematical units than there are programs in the software component library: a testament to the power of mathematics in building useful abstractions.

The next three subsections concentrate on major components depicted in the verifying compiler architecture. Verification condition generation is the topic of Section 3, with the automated prover described in Section 4 and the proof checker in Section 5.

3 Automated Generation of Verification Conditions

A key element of the verifying compiler is its VC (verification condition) generator, which takes relevant specifications and an implementation and produces a set of mathematical assertions that, if proven, guarantee the correctness of the

code. The proof rules that underlie VC generation and claims of soundness and completeness of those rules are discussed elsewhere [12][21][22][33][44]. At present we are building and experimenting with two versions of the VC generator, one that begins with the goals of the implementation and proceeds to derive VC's one at a time by sweeping backward through the code, and another that leads to a table of known assertions and assertions to be proved and that is effectively neutral with respect to the order in which VCs are generated. The nature of these tools is such that they will allow us to study the impact of inherent differences in the structure of VCs and the impact of different (but logically equivalent) styles of specification writing from an automation perspective. The VC generators are conceptually plug-compatible, meaning one can be used in the place of the other within the toolset architecture. Both share the structural and semantic framework common to RESOLVE [43] (though they differ slightly in terms of the language syntax they accept), so the comments on specification, design, and implementation mentioned in context in Section 4.1 (e.g., modularity issues) and Section 4.2 (e.g., swapping) are applicable to both approaches. The significant differences are in the proof rules underlying VC generation and the structure of VCs.

3.1 Goal-Directed VC Generation

The formal proof rules underlying the goal-directed approach are discussed in [21]. To illustrate the VCs generated by this approach, we have chosen benchmark #3 from [48]. The goal of this benchmark is to specify and implement a piece of code involving a user-defined abstract data type (ADT) *Queue* instead of code that uses structures such as integers or arrays which are usually built into languages. Specifically, this benchmark requires verification of the (iterative or recursive) code for a *Queue* reversal operation. In this section, we explain how our VC generator addresses this challenge in a modular fashion, using only the specification of *Queue*, without regard to how a *Queue* is represented and realized. This means that, among other advantages, the proof does not deal with nodes, pointers, and other possibly quirky details of *Queue* representations [36][50]. Those internal details will be verified when the data abstractions themselves are verified. This approach makes reasoning about pointers a relatively-independent verification problem as opposed to an intertwined problem that must be addressed in specification and verification of all object-based software [35].

Figure 5 contains the specification of a *Queue* flipping operation. This specification enhances, or extends, the specification of *Queue_Template* (given in an Appendix), where *Queues* are conceptualized mathematically as strings of entries by importing *String_Theory*. Details of the specification notation may be found in [34]. In general, concepts, such as *Queue_Template* may import one or more (highly reusable) mathematical theories as necessary. The *Queue* type defined in *Queue_Template* is bounded, so the component is parameterized both by the type of entries in the *Queue* and a maximum length. All *Queues* are conceptually con-

strained to be within this length. *Queue_Template* provides a set of basic, or primary, operations to manipulate *Queue* objects, such as *Enqueue*, *Dequeue*, and *Length*, and it is designed to avoid the need to copy shallow references or deep representations [20][34]. In the ensures clause of *Flip*, $\#Q$ denotes the input value of *Q* and “reverse” is a mathematical string operator.

Figure 6 contains a recursive implementation of the *Flip* operation. A recursive procedure must be annotated with a **decreasing** clause that is type-checked to be an ordinal. This progress metric must be shown to decrease with each recursive call in order to show termination. Loops in iterative implementations must be annotated additionally with an invariant through a **maintaining** clause as illustrated in the next subsection.

```

Enhancement Flipping_Capability for Queue_Template;
  Operation Flip( updates Q: Queue);
    ensures Q = reverse(#Q);
end Flipping_Capability;

```

Figure 5: Specification of a Queue Flipping Capability

```

Realization Recursive_Flipping_Realiz for
  Flipping_Capability of Queue_Template;
  uses Std_Boolean_Fac;
  Procedure Flip( updates Q: Queue);
    decreasing |Q|;

    Var E: Entry;
    If (Length(Q) /= 0) then
      Dequeue(E, Q);
      Flip(Q);
      Enqueue(E, Q);
    end;
  end Flip;
end Recursive_Flipping_Realiz;

```

Figure 6: A Recursive Implementation

The goal-directed VC generator can produce outputs in multiple forms, including an Isabelle-friendly version (not shown). Four VCs that arise for the flipping implementation are shown in Figure 7. Each VC is an implication and it includes the assumptions that may be used to prove the obligation following the implication. The types of variables are shown at the top, where *Str* is the name given for strings given in *String_Theory*. In the VCs, “o” denotes string concatenation, $\langle E \rangle$ denotes a string containing one entry *E*, $|\alpha|$ denotes the length of a string α .

It is important to note that the VCs are purely mathematical assertions. Any program variable (e.g., a *Queue Q*) that appears in a VC stands for its mathematical value (a mathematical string).

In the VCs, variable names with a preceding “?” (e.g., *?E*, *?Q*) along with *P_Val* are generated in the verification process. The assumptions in the VCs arise from constraints on objects (e.g., the constraints $|Q| \leq \text{Max_Length}$) and knowledge about the states based on the specifications of called operations or branch conditions.

```

Free Variables: Max_Length:*Z, min_int:*Z, max_int:*Z,
                  Q:*Str(*Entry), P_val:*N, ?E:*Entry,
                  ?Q:*Str(*Entry), E:*Entry

(((min_int <= 0) and (0 < max_int)) and
(Max_Length > 0)) and
((|Q| <= Max_Length) and (P_val = |Q| and
(|Q| /= 0 and Q = (<?E> o ?Q))))
=====>
(|?Q| < |Q|)

(((min_int <= 0) and (0 < max_int)) and
(Max_Length > 0)) and
((|Q| <= Max_Length) and (P_val = |Q| and
(|Q| /= 0 and Q = (<?E> o ?Q))))
=====>
(|reverse(?Q)| < Max_Length)

(((min_int <= 0) and (0 < max_int)) and
(Max_Length > 0)) and
((|Q| <= Max_Length) and (P_val = |Q| and
(|Q| /= 0 and Q = (<?E> o ?Q))))
=====>
(reverse(?Q) o <?E>) = reverse(Q)

(((min_int <= 0) and (0 < max_int)) and
(Max_Length > 0)) and ((|Q| <= Max_Length) and
(P_val = |Q| and |Q| = 0))
=====>
Q = reverse(Q)

```

Figure 7: VCs Generated for Recursive Flipping Realization

The goal-directed approach attempts to minimize both the number of VCs that are generated and the number of new variables that are introduced, and does other

simplification to ease the task of the back-end prover. To prove the correctness of the code the following VCs are generated and proved. Figure 7 shows only four VCs because the fifth one has been simplified, proved, and eliminated by the VC generator itself.

- Termination (VC 1)
- Requires clause of called operation *Dequeue* (simplified to true by VC generator)
- Requires clause of called operation *Enqueue* (VC 2)
- Ensures clause of *Flip* when if condition is true (VC 3)
- Ensures clause of *Flip* when if condition is false (VC 4)

The VCs in the figure took a fraction of a second to generate and 0.7 seconds for Isabelle to prove. In addition to proving similar conditions, for code involving loops, VCs to prove the validity of the (programmer-supplied) loop invariant are also generated [21].

3.2 Tabular VC Generation

In this second approach, VCs are generated into a locally defined XML format using a tool we have developed that implements the proof rules described formally in [22] and informally in [45]. The XML serves as an intermediate form that is translated into Unicode for human consumption, into input for Isabelle for automated proof, and into input for our own prover SplitDecision (see below) that incorporates special-purpose decision procedures for fragments of mathematical theories that are widely used in RESOLVE, e.g., strings, finite sets, and integers.

As outlined in [48], the tabular method generates one VC for each state in the program where the next statement involves either establishing a pre-condition for the next operation invoked, or a loop invariant or progress metric, or the post-condition for the operation being verified. The VC generator works in two phases. Rote VC generation introduces in the first phase a large number of mathematical variables: one for each program variable in each program state (between consecutive statements). For instance, x_i stands for the value of program variable x in state i ; in a program with n variables and m statements, the VCs constructed this way therefore collectively contain a total of approximately nm distinct mathematical variables at the end of phase one.

The simplification phase of the VC generator then applies a few theory-independent logical restructuring rules to each VC. The most obvious and useful simplification is simply substitution. Many of the nm variables are removed by this step. For instance, the hypotheses in the original VCs often contain clauses such as “ $x_{10} = x_9$ ”, with a substitution removing one of these two variables (and this clause). Any effective back-end automated prover would certainly be able to

apply such substitutions to achieve the same effect, but the VC generator does this before handing the VCs to an automated prover so the human-readable output of the VC generator is more concise. Some other work done in the second phase is intended to help the back-end prover by making structural changes that are expected to be fruitful because of the nature of the proof rules. An example is the use of automatic case analysis. In our experience, case analysis can only rarely be performed by a back-end automated prover: the requirement for human advice about appropriate cases seems to be the norm. The VC generator therefore strives to make it unnecessary: the simplification phase divides each VC into appropriate cases based on the path structure of the code from which the VCs arise so the automated prover need not be asked to infer information that was previously known but effectively obfuscated by the initial phase of VC generation.

The net result is that, in the tabular approach, there are normally many VCs even for a relatively short piece of code; On the other hand, each VC tends to be relatively small and apparently “prover-friendly”. Experience to date shows that the vast majority of such VCs can be discharged without difficulty by both the automated provers we are using—based on fairly limited mathematical theory libraries (in the case of Isabelle) and limited built-in mathematical knowledge about the theories involved in the specifications (in the case of SplitDecision).

Here is an example of how the process works. Consider an enhancement that allows a client program to form unions and intersections of two sets given the *Set-Template* contract (see Appendix B). The specification of this enhancement is given in Figure 8 and an implementation is shown in Figure 9.

An important feature of this code is the use of the swap statement (“:=” operator) to swap, or exchange, the values of two variables of the same type. Swapping is a staple of RESOLVE [20] and it replaces the typical variable-to-variable assignment in other languages. Among other things, swapping admits efficient implementation for every type without introducing the need for reference semantics and the associated problems of aliasing that otherwise would significantly complicate modular verification [49]. In the code in Figure 9, swapping is used for two purposes. First, the programmer makes sure that the iteration is over the smaller of the two sets provided as arguments (for efficiency). Then the local *Set* variable *tmp* is used to hold those elements of the incoming *t* that are also in the incoming *s*, i.e., those that are in their intersection. At the end of the loop, since the value of *tmp* is the intersection of the original arguments and this value is supposed to be returned in *t*, these two variables are swapped. This iteration idiom is common in RESOLVE programs that use collection ADTs [20][50].

```

contract SetUniteAndIntersect enhances SetTemplate
  procedure UniteAndIntersect (updates s: Set,
                               updates t: Set)
    ensures
      s = #s union #t and t = #s intersection #t
  end SetUniteAndIntersect;

```

Figure 8: An Operation to Union and Intersect Sets

```

realization Iterative2 implements SetUniteAndIntersect

  procedure UniteAndIntersect (updates s: Set,
                               updates t: Set)

    variable ss, ts: Integer
    variable tmp: Set
    ss := Size (s)
    ts := Size (t)
    if IsGreater (ts, ss) then
      s :=: t
    end if
    loop
      maintains s union t = #s union #t and
        (s intersection t) union tmp =
          (#s intersection #t) union #tmp and
          t intersection tmp = {}
      decreases |t|
    while not IsEmpty (t) do
      variable x: Item
      RemoveAny (t, x)
      if not IsMember (s, x) then
        Add (s, x)
      else
        Add (tmp, x)
      end if
    end loop
    t :=: tmp
  end UniteAndIntersect

end Iterative2

```

Figure 9: An Implementation to Unite and Intersect Sets

Our VC generator produces 38 VCs for this code. They range in apparent complexity through the spectrum of samples shown in Figure 10, which are cut directly from the human-readable output produced by the tabular VC generator. A few VCs that look like #15 are evidently the most difficult of these to prove, but even in this case the conclusion follows from only the hypotheses 2 and 7 and a couple general-purpose properties of finite sets. VCs near the end of the code are hardly more difficult than those near the beginning. The conclusion of VC #5 follows from the property that the empty set is a right identity for intersection; there are a number of other VCs like this one. VC #10 is typical of a few more VCs: its conclusion is simply one of the hypotheses. The conclusion of VC #35, like a few other VCs of the same nature, follows from hypothesis 1 and a couple other general-purpose properties of finite sets.

Both Isabelle and SplitDecision prove all 38 VCs from this code automatically. Isabelle does it by expanding union, intersection, and difference into their definitions based on set membership and then relying on its considerable power to simplify boolean expressions. SplitDecision uses a similar approach encoded as a set of structural rewrites (in this case, involving expressions with finite-set operators such as union and cardinality) that simplify each VC as far as possible—and all these VCs are simplified to “true”.

Our web interface to the toolset, and the VC generator as well as both back-end automated provers, currently run on a single-core 2.8 GHz Intel Xeon processor with 1 GB of memory; Redhat Enterprise Linux 4 is the operating system, and Apache Tomcat is the web server. The response time on clicking on the “Verify” button for this code results is quite reasonable: VC generation takes a small fraction of a second, and Isabelle proves all 38 VCs in about 2.7 seconds while SplitDecision proves them all in about 2.0 seconds of wall-clock time.

To be sure, not all VCs involving strings, finite sets, and integers are automatically discharged for all the examples we have tried. We currently let Isabelle time-out at 2 seconds on each VC, and occasionally it does so without having completed a proof or refutation. Either Isabelle or SplitDecision might also simplify a VC to something else with no further progress in sight. Considerable empirical evidence will need to be obtained in order to generalize comfortably about why some VCs turn out to be difficult. Some known challenges for the future include VCs involving permutations of strings (which arise in the specifications and annotations of sorting algorithms), those using a function that maps the entries of a string to the finite set containing those entries, and situations where the proof of a VC involves an algebraic property whose importance was not apparent at the time of theory development (e.g., that concatenation of strings is commutative in arguments to the *is_permutation* predicate). We are confident that these sorts of VCs can be handled effectively with additional library development for Isabelle and additional rules for SplitDecision.

Verification Condition #5 (state index: 7, loop invariant)	
1:	$l_{tol} > l_{sol}$
	$\Rightarrow s_o \cap \emptyset = \emptyset$
...	
Verification Condition #10 (state index: 9, requires clause)	
1:	$t_s \neq \emptyset$
2:	$\wedge s_s \cup t_s = s_o \cup t_o$
3:	$\wedge (s_s \cap t_s) \cup tmp_s = (s_o \cap t_o) \cup \emptyset$
4:	$\wedge t_s \cap tmp_s = \emptyset$
5:	$\wedge l_{tsl} \geq 0$
6:	$\wedge is_initial(x_9)$
7:	$\wedge l_{tol} \leq l_{sol}$
	$\Rightarrow t_s \neq \emptyset$
...	
Verification Condition #15 (state index: 15, loop invariant)	
1:	$t_s \neq \emptyset$
2:	$\wedge s_s \cup t_s = t_o \cup s_o$
3:	$\wedge (s_s \cap t_s) \cup tmp_s = (t_o \cap s_o) \cup \emptyset$
4:	$\wedge t_s \cap tmp_s = \emptyset$
5:	$\wedge l_{tsl} \geq 0$
6:	$\wedge is_initial(x_9)$
7:	$\wedge x_{10} \in t_s$
8:	$\wedge is_initial(x_{12})$
9:	$\wedge l_{tol} > l_{sol}$
10:	$\wedge x_{10} \notin s_s$
	$\Rightarrow (s_s \cup \{x_{10}\}) \cup (t_s \setminus \{x_{10}\}) = t_o \cup s_o$
...	
Verification Condition #35 (state index: 17, ensures clause)	
1:	$s_{16} \cup \emptyset = t_o \cup s_o$
2:	$\wedge (s_{16} \cap \emptyset) \cup tmp_{16} = (t_o \cap s_o) \cup \emptyset$
3:	$\wedge \emptyset \cap tmp_{16} = \emptyset$
4:	$\wedge l_{tol} > l_{sol}$
	$\Rightarrow s_{16} = s_o \cup t_o$

Figure 10: VCs for Unite and Intersect Set Implementation

Neither the goal-directed approach nor the tabular approach currently generates VCs for all of the features the RESOLVE language. However, no known technical barriers remain for VC generation. For example, one important next step is to add VC generation for components that include new data type representations (e.g., a realization of *SetTemplate*). The idea is that the realization is coded with annotations similar to those in Hoare’s framework for proof of data representation [24][44](i.e., with a representation invariant and an abstraction relation for each new type). This results in proof obligations for the implementations of the component’s operations.

There are numerous issues that can make it difficult to prove the generated VCs, and they need to be addressed. They concern mathematical theories, their use in specifications, and programming paradigms. One challenge for math specialists is suitable theory development (in terms of both definitions and theorems presenting algebraic properties about them) that is necessary to minimize the usage of both existential and universal quantifiers usage. For example, though auxiliary variables can be used in RESOLVE to automate verification in the presence of existential quantifiers (e.g., *Swap_First_Entry* operation in *Queue_Template* specification in the Appendix) following the lines of [31], suitable string theory development (e.g., with the use of a string “part between” definition and supporting theorems) can eliminate the need for quantifiers to a large extent [32]. The challenge for specifiers is devising specifications that facilitate ease of mechanization, yet are easy for people to understand. Just as programmers are taught to write “good” code, software developers should learn to write “good” specifications. The challenge for language designers is in taking into account ease of verifiability considerations in addition to efficient performance. Finally the tools should provide meaningful feedback to programmers when there are errors, and assist them in the process of eventually developing verified software.

4 Automated Proving of VCs

This section summarizes two complementary approaches for proving generated verification conditions. The first section explains how we import our theory developments to Isabelle and illustrates what can and cannot be proved using only “auto” mode. The second subsection discusses a custom-built tool for proving RESOLVE VCs, called SplitDecision.

4.1 Automated Proving Using Isabelle

A natural question about proving VCs using a particular theorem proving tool such as Isabelle is how the VCs should be expressed. We take the approach of

creating VCs that are consistent with RESOLVE mathematical theories instead of generating VCs directly in the language of the prover (Isabelle) that use prover-specific theories provided by other Isabelle users or its developers. This decision allows us to remain decoupled from Isabelle: we import RESOLVE mathematical theories for use in the VC proofs by encoding them for Isabelle. There are many other choices for an off-the-shelf automated prover, and while Isabelle has many merits, it is not necessarily the ultimate back-end automated prover for RESOLVE VCs. For the examples used in this paper, we have created both a RESOLVE finite set theory and a RESOLVE string theory within Isabelle. Both theories use types defined in Isabelle as witnesses to the consistency of the theories. However, once these off-line existence proof obligations are dispatched, only lemmas defined in the RESOLVE theories are used by Isabelle to prove VCs.

This approach both reduces the number of errors in the importing process (by providing witnesses to ensure consistency) and ensures that the theorem proving tool used “understands” the mathematics in the VCs generated. We also note that the particular method of VC generation (goal directed or tabular) is unimportant to the proof of the VC; both sources of VCs result in assertions that can be processed and proved by Isabelle.

The first examples presented in this paper have specifications in terms of mathematical string theory. These mathematical strings are used as a mathematical model of the behavior of the *Queue* ADT, which is similar to what can be done using sequences in JML [37]. However specifications and code are two distinct entities in RESOLVE; specifications are statements in mathematics and are not intended to be executable.

Informally, strings over a given type *obj* are intended to have a model that is exactly the elements of obj^* , where $*$ is the Kleene star. However, a theory such as the one in Figure 10 is self-contained and might have other models; that is, a theory is defined not by exhibiting a particular model, but rather by its axioms. Of course, it is important to know that such a model exists and that we are not describing an inconsistent theory. A few functions are defined inductively for strings: concatenation, length, and reverse. In Figure 10, we state some simple lemmas without proofs. The proofs of these lemmas follow easily (but not automatically) from the axioms and definitions. Once proved with human assistance, though, these lemmas can be used in proofs of VCs, exactly like the axioms and definitions.

String Type Signature

$$\begin{aligned} \text{string} &\stackrel{\text{def}}{=} \text{string}(\text{obj}) \\ \Lambda &: \text{string} \\ \text{ext} &: \text{string} \times \text{obj} \longrightarrow \text{string} \end{aligned}$$
String Axioms

1. $\text{ext}(s, x) \neq \Lambda$
2. $\text{ext}(s_1, x_1) = \text{ext}(s_2, x_2) \Rightarrow s_1 = s_2 \wedge x_1 = x_2$
3. $\forall S \in \mathcal{P}(\text{string}) : (\Lambda \in S \wedge \forall x, s : (s \in S \Rightarrow \text{ext}(s, x) \in S)) \Rightarrow S = \text{string}$

Function Definitions

1. $\langle _ \rangle : \text{obj} \longrightarrow \text{string} \stackrel{\text{def}}{=} \langle x \rangle = \text{ext}(\Lambda, x)$
2. $|_| : \text{string} \longrightarrow \mathbb{N} \stackrel{\text{def}}{=} |\Lambda| = 0 \wedge |\text{ext}(s, x)| = |s| + 1$
3. $* : \text{string} \times \text{string} \longrightarrow \text{string} \stackrel{\text{def}}{=} (s * \Lambda = s) \wedge (s_1 * \text{ext}(s_2, x) = \text{ext}(s_1 * s_2, x))$
4. $\text{reverse} : \text{string} \longrightarrow \text{string} \stackrel{\text{def}}{=} \text{reverse}(\Lambda) = \Lambda \wedge \text{reverse}(\text{ext}(s, x)) = \langle x \rangle * \text{reverse}(s)$

Useful Lemmas

1. lemma EmptyNotSingle: $\Lambda \neq \langle x \rangle$
2. lemma IdofEmpty : $\Lambda * \alpha = \alpha$
3. lemma LenofSingle : $|\langle x \rangle| = 1$
4. lemma LenofCat : $|\alpha * \beta| = |\alpha| + |\beta|$
5. lemma AssocCat : $\alpha * (\beta * \gamma) = (\alpha * \beta) * \gamma$
6. lemma ReverseofReverse : $\text{reverse}(\text{reverse}(\alpha)) = \alpha$
7. lemma ReverseofCat : $\text{reverse}(\alpha * \beta) = \text{reverse}(\beta) * \text{reverse}(\alpha)$
8. lemma LenofReverse: $|\text{reverse}(\alpha)| = |\alpha|$

Figure 10: String Theory

First we consider the VCs in Figure 7 that are generated from the code to flip a queue in Section 3.1. The first VC relates to showing termination by showing that the length of the *Queue* passed to the recursive call is strictly smaller than the length of the original *Queue*. Isabelle can prove this easily based on lemmas about the operators involving strings.

Verification Condition #1 from Figure 7 (recursion termination)

$$\begin{aligned}
& \text{min_int} \leq 0 \\
\wedge & 0 < \text{max_int} \\
\wedge & \text{Max_Length} > 0 \\
\wedge & |Q| \leq \text{Max_Length} \\
\wedge & P_val = |Q| \\
\wedge & |Q| \neq 0; \\
\wedge & Q = \langle E1 \rangle * Q1
\end{aligned}$$

$$|Q1| < |Q|$$

The second VC corresponds to showing that the ensures clause is true when the if condition holds, i.e., *Queue* length is not 0. It is proven easily by Isabelle based on algebraic properties of reverse and concatenation that are captured in a few oft-used lemmas.

Verification Condition #3 from Figure 7 (ensures clause for recursive case)

$$\begin{aligned}
& \text{min_int} \leq 0 \\
\wedge & 0 < \text{max_int} \\
\wedge & \text{Max_Length} > 0 \\
\wedge & |Q| \leq \text{Max_Length} \\
\wedge & P_val = |Q| \\
\wedge & |Q| \neq 0 \\
\wedge & Q = \langle E1 \rangle * Q1
\end{aligned}$$

$$\text{reverse}(Q1) * \langle E1 \rangle = \text{reverse}(Q)$$

Next, we consider proof of VCs arising from the *Set* example in Section 3.2. The finite set theory used by RESOLVE provides all of the expected set functions and predicates, such as set membership, union, intersection, set difference, singleton, and subset. It also provides the set cardinality function that represents the size of the set. We note that, without the cardinality function, this theory is decidable [15]. Moreover, the cardinality function also restricts the combinations of decision procedures for other theories (for example, string theory above) because of the potential for “crosstalk” involving the use of natural numbers as the range of the length function.

We use two main rules to simplify formulas that involve cardinality: $|a \cup b| = |a| + |b| - |a \cap b|$ and $|a - b| = |a| - |a \cap b|$. Of course the cardinality of a singleton is one and the cardinality of the empty set is 0. . These rules have been powerful enough to handle the VCs generated by examples similar to the one in Section 3.2 but of course may need to be extended to handle VCs arising from other code in the future. As noted earlier, for technical reasons we do not expect the process of adding new lemmas to Isabelle’s theories or new rules to SplitDecision to end with the ability to prove all true VCs. We do, however, expect the frequency with which the math specialist needs to make such improvements to decrease conti-

nually as more improvements are made for a particular theory.

As previously noted, for the examples used in the tabular method of generating VCs which involve sets, Isabelle is able to prove all of the VCs. Some involve simple manipulations of the lengths of finite sets while others require more involved simplifications using the definitions of the set functions and predicates. We present two VCs, along with what the proofs entail.

The first VC exemplifies the interactions between cardinality of finite sets and the definitions of sets, and is shown below. The VC corresponds to proving that the loop body as shown in Figure 9 actually does decrease the progress metric. The proof, performed automatically by Isabelle via the “auto” tactic, follows by applying various simplification rules that first simplify the consequent into $|t\delta| - |t\delta \cap \{x_{10}\}| < |t\delta|$. Then, an additional rule which simplifies intersections if one of the arguments is a subset of the other argument is used to further simplify the formula to $|t\delta| - |\{x_{10}\}| < |t\delta|$. Finally, this formula can be shown to be “true” using algebra and the fact that $|\{x_{10}\}| = 1$. The proof requires some bookkeeping and use of simple lemmas.

Verification Condition #31 (state index: 15, loop metric)

$$\begin{aligned}
 & t_8 \neq \emptyset \\
 & \wedge s_8 \cup t_8 = t_0 \cup s_0 \\
 & \wedge (s_8 \cap t_8) \cup tmp_8 = (t_0 \cap s_0) \cup \emptyset \\
 & \wedge t_8 \cap tmp_8 = \emptyset \wedge |t_8| \geq 0 \\
 & \wedge is_initial(x_9) \\
 & \wedge x_{10} \in t_8 \\
 & \wedge is_initial(x_{12}) \\
 & \wedge |t_0| > |s_0| \\
 & \wedge x_{10} \notin s_8
 \end{aligned}$$

$$|t_8 \setminus \{x_{10}\}| < |t_8|$$

The VC integral to the correctness of the computation is shown next. This VC involves proving that a clause of the loop invariant holds at the end of the loop. In the proof of this VC, instead of using specific facts about cardinality, the proof involves expanding the extensionality axiom, along with using the definitions of set membership with respect to the union, intersection and set difference functions. For example, an item is in the union of two sets if it is an element of either set. Once the definitions are expanded, it becomes a matter of propositional and (simple) predicate logic manipulations that are performed easily by Isabelle.

Verification Condition #22 (state index: 15, loop invariant)

$$\begin{aligned}
& t_8 \neq \emptyset \\
& \wedge s_8 \cup t_8 = s_0 \cup t_0 \\
& \wedge (s_8 \cap t_8) \cup tmp_8 = (s_0 \cap t_0) \cup \emptyset \\
& \wedge t_8 \cap tmp_8 = \emptyset \\
& \wedge |t_8| \geq 0 \\
& \wedge is_initial(x_9) \\
& \wedge x_{10} \in t_8 \\
& \wedge is_initial(x_{14}) \\
& \wedge |t_0| \leq |s_0| \\
& \wedge x_{10} \in s_8 \\
\hline
& (s_8 \cap (t_8 \setminus \{x_{10}\})) \cup (tmp_8 \cup \{x_{10}\}) = (s_0 \cap t_0) \cup \emptyset
\end{aligned}$$

The current theories available for use by Isabelle to prove VCs resulting from RESOLVE programs include finite sets and string theory. Work proceeds on adding more RESOLVE mathematical theories such as tree theory and graph theory to Isabelle. More concrete examples of possible challenges include verifying sorting implementations. These implementations make use of mathematical functions that abstract the notion of sorting, e.g., the statement that every entry in one string is less than or equal to any entry in another string (where “less than or equal to” is a proxy for any total pre-ordering on the actual entry type). These highly generic definitions, while necessary to deal with fully generic code of typical RESOLVE software components, can complicate automated proofs. Furthermore, identification of ways to streamline, augment, and check how RESOLVE mathematical theories are imported to the provers is also ongoing. These theories were developed with ease of specification in mind. In order to support automated verification as well, extra lemmas may be necessary. Moreover, there may be some degree of redundancy in the theories that can be reduced both to aid the specifiers and to aid the automated proof tools.

4.2 Automated Proving with SplitDecision

When not using its 'auto' mode, Isabelle is an interactive proof assistant. Often a proof assistant will halt its execution and present the user with alternatives for how to proceed, in an effort to reduce the number of “blind alleys” that it might follow in its proof efforts. This dependence on outside intervention is not compatible with the desire to have VCs proven completely automatically. Therefore, in parallel with our exploration of Isabelle's strength in 'auto' mode, we are also pursuing the development of a new tool geared specifically for push-button RESOLVE VC proving called SplitDecision.

SplitDecision leverages domain-specific mathematical knowledge along with general-purpose strategies for reducing logical formulae in order to simplify RESOLVE VCs while maintaining logical equivalence. Each action it takes makes progress toward the ultimate goal of a “true” or “false” conclusion, and it never backtracks. It continues to work unidirectionally until truth or falsehood is established, or until no more simplifications are applicable.

The simplifications SplitDecision is allowed to perform are defined by specialized decision procedures tuned to handle some of the most-used mathematical theories involved in RESOLVE specifications. The goal of this work is for SplitDecision to be able to prove/disprove any RESOLVE VC, or at least reduce it to a form that a proof assistant will need no further guidance to prove/disprove. So far, decision procedures for Presburger arithmetic and a specific subclass of assertions in mathematical string theory have been implemented. Recent theoretical results [3] indicate that a theory of arrays can be handled similarly, and research by both computer scientists and logicians continues in an attempt to address more entities such as sets and trees.

SplitDecision is an ongoing project. It was originally built with the intent of merely making the VCs simpler for Isabelle. It has, instead, grown to be able to verify outright a large fraction of the VCs we have examined thus far. At the time of writing, a full decision procedure for strings with certain technical restrictions is implemented, and a substantial amount of progress has been made on such a procedure for finite sets.

5 Proof Checking

While the VCs that arise from the software specialist's program are to be dispatched automatically, this process is supported by theorems defined by the math specialist. As has already been discussed, a “false theorem” would compromise the soundness of the entire system and as a result it is important that any putative theorems themselves be verified. However, since these theorems might be of arbitrary complexity, in general, they cannot be verified automatically. To this end, the math specialist may provide explicit proofs in a formalized language that can then be checked by the system. One such proof checking system integrated within the RESOLVE framework is discussed in detail in [46]. Since these theories are organized into libraries around mathematical concepts (e.g., strings) which are reused by many different programming objects (e.g., stacks, queues, and lists), the overhead associated with writing and checking proofs of non-trivial theorems is amortized over time and reuse.

For example, consider the theorem $\forall \alpha, \beta, (reverse(\alpha \bullet \beta) = reverse(\beta) \bullet reverse(\alpha))$ from Section 2. This theorem was necessary to complete the proof of the third VC in Figure 7, which arose from the verification of the Queue Flipping implementation. It is found in a RESOLVE mathematical unit for String_Theory called a *précis* that contains a summary of the theorems in a theory. Here is what the theorem might look like in context:

```

Precis String_Theory;
...
  Theorem S1: For_all alpha, beta:
    reverse(alpha o beta) =
      reverse(beta) o reverse(alpha);
...
end String_Theory;

```

The *précis* is an interface to mathematical theory developments and it provides only theorems, such as the one above, so that they can be used by clients such as the automated prover, the math specialist, or the software specialist. Corresponding proofs are provided in a separate *proof unit*, also a part of the RESOLVE framework. A design goal of the RESOLVE proof-writing language is to highlight the separation between results, which are used by consumers like the prover, and corresponding proofs, which are used only to ensure soundness. We note that math specialists are free to use other tools, such as automated provers or proof assistants, to help generate the required proofs.

One primary design challenge in developing a proof checker is balancing the needs of an explicit, formal language with the expectations of trained mathematicians rather than programmers. The current version of the RESOLVE proof checker is described in [46] with an illustrative example.

6 Related Work and Conclusions

There are many forks in the road of research toward verified software [38] and there are many fundamental questions [29]. Considerable related work lies down the paths *not* taken by the project team, all of which offer hope for progress that other researchers will and should continue to pursue. The rationale for each major decision made so far in our efforts is summarized below, along with relevant related work.

First, it is important to establish the scope of software to be considered. There are a number of options that fall along a spectrum including (a) concurrent/distributed software with a focus on a limited set of properties to be verified, and (b) full behavioral verification of sequential software. Ultimately, the vision is to be able to specify and verify all software, but this project focuses on the latter type—specifically, on sequential software *components*, because all useful software today is component-based. The work is complementary in principle to most recent work on finite-state model-checking, etc. (which suffers equally from the technical drawback that it cannot address the entire problem, e.g., software systems whose appropriate models are manifestly not finite-state).

Within the context of software components, a major choice involves the goal—indeed, the definition—of “verification”. The question is whether to prove that software satisfies (a) some fixed set of properties, e.g., that there are no null dereferences or that no array indices are out of range [10]; or (b) a partial specification of behavior, e.g., that adding an element to a queue increases the length of the queue by one (without stipulating anything about, for example, the order of the actual queue elements) [40]; or (c) a complete behavioral specification of what the software is supposed to do and what it is not supposed to do [1][28][37][43][48][51]. Our goal is the last one. This “full functional” notion of verification is certainly the most difficult to achieve, but accounting for complete behavior will be required eventually, and this approach subsumes the others.

Another important decision is whether to try to verify software written in (a) some industrial language, e.g., Java or C or C++ or C#, or (b) a clean language such as RESOLVE [43] that lends itself to systematic exploration of fundamental research questions related to verified software. A critical point is that RESOLVE has been designed to support *modular* verification based on the principle of design-by-contract [40]: each component is verified in isolation, once and for all, rather than being verified every time it is used in a client program. This property is necessary for the verified software paradigm to be scalable to practical-size programs. RESOLVE also is rich enough that software components written in it are at least as general and efficient as their counterparts in the above languages, and it has been used successfully to develop commercial software [27].

A final major decision involves mechanical theorem-proving tools used to prove VCs. There are many candidates for such a back-end. We currently use Isabelle, which is intended to be a proof assistant for professional mathematicians but also can serve as a fully automated prover. Isabelle offers several technical benefits as it serves now as a proxy for the ultimate automated prover. Like RESOLVE’s specification language, it is not limited to first-order logic, but instead supports higher-order logic (Isabelle/HOL [41]). Crucially, Isabelle provides facilities both to define new mathematical theories (with associated lemmas, theorems, and corollaries), and to incorporate new mathematical results into its automated reasoning. This ability to create new theories is integral to the task at hand because it provides the right layer of abstraction between the specifics of how Isabelle performs proofs and the mathematical theories used in RESOLVE specifications.

The rest of this section elaborates on most related tools. Like RESOLVE, Spec# focuses on specifying an imperative language and providing compile-time push-button verification. Unlike RESOLVE, Spec# is based on a preexisting commercial language, C#. It does not significantly alter that language to ease verification, though it does offer at least one new feature to ease specification: non-null references. Whereas RESOLVE compiles mathematical assertions directly to VCs, Spec# compiles first to an intermediate language, Boogie [4], intended to generalize the verification process. Spec# also implements dynamic checks in ad-

dition to static ones so that specified components may properly interact with unspecified ones.

The Coq IDE provides a graphical front-end for working with the Coq proof management system. The user is able to load up a Coq file and incrementally step through portions of the file, getting useful feedback about where and why problems are encountered. Like the RESOLVE system, the Coq IDE provides immediate, push button feedback on the veracity of mathematical statements with the additional feature of helping to establish where exactly the problem lies. Coq is not intended as a programming language but provides useful proving and proof-checking features for a verification system [9].

ESC/Java2 [13] is a system that applies the Extended Static Checking [10] system to the Java programming language. ESC/Java2 operates on JML notations within a Java program. Like the RESOLVE system, it strives to statically detect errors in code. Unlike RESOLVE, it does not strive to be sound or relatively complete, but simply to flag a class of potential problems. Like Spec#, ESC/Java2 builds on top of an existing commercial language, and encourages incremental specification rather than full program specification. When combined with a development environment, such as when used with a plugin for Eclipse [14], it provides limited push-button verification. Aside from ESC/Java2, there are a number of different verification efforts underway that also use JML to specify Java code [7]. One example is the work in [29] for translating Java/JML correctness checks in to a Hoare-style logic processed by Isabelle for recursive functions. Verification based on JML is also concerned with keeping up with evolving Java language [8].

Unlike JML and Spec#, some verification tools are based on subsets of popular languages. The focus of SPARK (based on Ada) is safety-critical software; current SPARK tools can verify some of the earlier benchmarks in [48].

The Verisoft system is unique in that it is intended to address the verification of an entire system down to the level of hardware and device drivers. For example, the group's work includes a developing a verified compiler for C0, a restricted subset of C [39].

Some tools are intended for more general use. For example, the Why/Krakatoa/Caduceus platform can produce VCs from annotated Java and C source code [17]. It employs an intermediate programming language, *Why* for the generation of VCs which can be output in the format of several different third party provers, including Isabelle and Coq.

Software verification is a grand challenge. To address the challenge, we need to automate solutions to a variety of subproblems. This paper illustrates the progress we have made by discussing in detail a toolset that we have developed for push-button software verification in RESOLVE, a language designed to be amenable for verification. But obstacles to automation remain. They range from fine-tuning techniques for mathematical theory and software development to even more effective techniques for automating proofs.

Acknowledgments

This grant is funded in part by U. S. National Science Foundation (NSF) grants CCF-0811748, CCF-0811737, DMS-0700174, DMS-0701187, DMS-0701260, and DUE-0633506 and a grant from the John Templeton Foundation. Our thanks are due to Bill Ogden for his numerous contributions to the foundations of this research.

Appendix A: Specification of Queue_Template

```

Concept Queue_Template( type Entry;
                        evaluates Max_Length: Integer);
uses Std_Integer_Fac, String_Theory;
requires Max_Length > 0;

Type Family Queue is modeled by Str(Entry);
exemplar Q;
constraint |Q| <= Max_Length;
initialization ensures Q = empty_string;

Operation Enqueue( alters E: Entry;
                   updates Q: Queue);
requires |Q| < Max_Length;
ensures Q = #Q o <#E>;

Operation Dequeue( replaces R: Entry;
                   updates Q: Queue);
requires |Q| /= 0;
ensures #Q = <R> o Q;

Operation Swap_First_Entry( updates E: Entry;
                             updates Q: Queue);
requires |Q| /= 0;
ensures there exists Rem: Str(Entry) such that
    #Q = <E> o Rem and Q = <#E> o Rem;

Operation Length(restores Q: Queue): Integer;
ensures Length = (|Q|);

Operation Rem_Capacity(restores Q: Queue): Integer;
ensures Rem_Capacity = (Max_Length - |Q|);

Operation Clear(clears Q: Queue);

end Queue_Template;

```

Appendix B: Specification of SetTemplate

```

contract SetTemplate (type Item)
  uses UnboundedIntegerFacility

  math subtype SET_MODEL is finite set of Item

  type Set is modeled by SET_MODEL
    exemplar s
    initialization ensures
      s = empty_set

  procedure Add (updates s: Set, clears x: Item)
    requires x is not in s
    ensures s = #s union {#x}

  procedure Remove (updates s: Set, restores x: Item,
    replaces x_copy: Item)
    requires x is in s
    ensures s = #s \ {x} and x_copy = x

  procedure RemoveAny (updates s: Set, replaces x: Item)
    requires s /= empty_set
    ensures x is in #s and s = #s \ {x}

  function IsMember (restores s: Set, restores x: Item):
    control
    ensures IsMember = (x is in s)

  function IsEmpty (restores s: Set): control
    ensures IsEmpty = (s = empty_set)

  function Size (restores s: Set): Integer
    ensures Size = |s|
end SetTemplate

```

References

- [1] Aichernig B K, Larson P G (1997) A proof obligation generator for VDM-SL, In: *FME 4, LNCS 1313*. Springer-Verlag.
- [2] Alkassar E, Hillebrand M, Leinenbach D, Schirmer N W, Starostin A (2008) The Verisoft Approach to System Verification. In: *Proc. VSTTE 2008*. Springer-Verlag.
- [3] Barnes J, *High-Integrity Ada: The Spark Approach*, Addison-Wesley, 2002.
- [4] Barnett M, Chang B E, et al (2006) Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In: de Boer F S, Bonsangue M M, and de Roever W-P (eds.) *Fourth International Symposium on Formal Methods for Components and Objects*, LNCS 4111. Springer, pp 364-387.
- [5] Barnett M, Leino K R, Schulte W (2004) The Spec# Programming System: An Overview. In: Burdy L and Huisman M (eds.) *Construction and Analysis of Safe, Secure and Interoperable Smart devices International Workshop*, LNCS 3362. Springer, pp 49-69.
- [6] Bradley A, Manna Z (2007) *The Calculus of Computation: Decision Procedures with Applications to Verification*. Springer.
- [7] Burdy L, Cheon Y, Cok D, Ernst M, Kiniry J, Leavens G T, Leino K R M, Poll E (2005) An overview of JML tools and applications. *STTT* 7(3):212–232.
- [8] Cok D (2008) Adapting JML to generic types and Java 1.6. In: *Proc. SAVCBS*, pp 27-35.
- [9] The Coq Proof Assistant Reference Manual Version v8.1. <ftp://ftp.inria.fr/INRIA/coq/current/doc/Reference-Manual.pdf>.
- [10] Detlefs D L, Leino K R M, Nelson G, Saxe, J B (1998) Extended Static Checking. *Research Report 159*, Compaq Systems Research Center.
- [11] DeMillo R A, Lipton R J, Perlis, A J (1977) Social processes and proofs of theorems and programs. In: *Proc. POPL*, ACM, pp 206-214.
- [12] Ernst G W, Hookway R J, Ogden W F (1994) Modular Verification of Data Abstractions with Shared Realizations. *IEEE TSE* 20(4):288-207.

- [13] ESC/Java2. <http://secure.ucd.ie/products/opensource/ESCJava2/>. Accessed December, 2008.
- [14] EscJava2 Eclipse Plug-in. <http://kind.ucd.ie/products/opensource/ESCJava2/escjava-eclipse/updates/>. Accessed December, 2008.
- [15] Ferro, A., E. G. Omodeo and J. T. Schwartz, Decision procedures for elementary sublanguages of set theory. I. Multi-level syllogistic and some extensions, *Communications on Pure and Applied Mathematics* 33 (1980), pp. 599–608.
- [16] Fetzner J A (1988) Program verification: the very idea. *CACM* 31:1048-1063.
- [17] Filiâtre J, Marché C (2007) The Why/Krakatoa/Caduceus platform for deductive program verification. In: Damm W and Hermanns H (eds.), *Proc. CAV 19*, LNCS 4590. Springer-Verlag.
- [18] Ghilardi S, Nicolini E, Ranise S, Zucchelli D (2007) Decision Procedures for Extensions of the Theory of Arrays. *Annals of Mathematics and Artificial Intelligence* 50(3-4):231-254.
- [19] Harrison J (2008) Formal proof—theory and practice. *Notices of the AMS* 55:1395-1406.
- [20] Harms D E and Weide B W (1991) Copying and swapping: Influences on the design of reusable software components. *IEEE TSE* 17(5):424–435.
- [21] Harton H K, Sitaraman M, Krone J (2008) Formal Program Verification. In: Wah B (ed.), *Wiley Encyclopedia of Computer Science and Engineering*. John Wiley & Sons.
- [22] Heym, W (1995) *Computer Program Verification: Improvements for Human Reasoning*. Ph.D. thesis, The Ohio State University, 1995.
- [23] Hoare C A R (1969) An axiomatic basis for computer programming. *CACM* 12:576-585.
- [24] Hoare C A R (1972) Proof of Correctness of Data Representations. *Acta Informatica* 1, 271-281

- [25] Hoare C A R (2003) The verifying compiler: a grand challenge for computing research. *JACM* 50:63-69.
- [26] Hoare C A R, Misra J, Leavens G T, Shankar N (2007) *The Verified Software Initiative: A Manifesto*, April 2007, <http://qpq.csl.sri.com/vsr/manifesto.pdf/view>. Accessed December 2008.
- [27] Hollingsworth J E, Blankenship L, Weide B W (2000) Experience report: using RESOLVE/C++ for commercial software. In: Schneir B (ed.) *Proc. FSE*. ACM, pp 11-19.
- [28] Jackson D (2006) *Software Abstractions: Logic, Language, and Analysis*. MIT Press.
- [29] Jones C B (2003) The early search for tractable ways of reasoning about programs. IEEE, *Annals of the History of Computing*, 25(2):26-49.
- [30] Karabotsos G, Chalin P, James P R, Giannas L (2008) Total correctness of recursive functions using JML4 FSPV. In: *Proc. SAVCBS*, pp 19-27.
- [31] King J C (1970) *A Program Verifier*. Ph.D. dissertation, Carnegie Tech, 261 pp.
- [32] Kirschenbaum J, Harton H K, Sitaraman M (2008), A Case Study in Automated Verification. In: Shankar N (ed.) *Proc. CAV Workshop on Automated Formal Methods*.
- [33] Krone J (1988) *The Role of Verification in Software Reusability*. PhD. Thesis, The Ohio State University..
- [34] Kulczycki G, Sitaraman M, Yasmin N, Roche K (2008) *Formal Specification*. In: Wah B (ed.), *Wiley Encyclopedia of Computer Science and Engineering*. John Wiley & Sons.
- [35] Kulczycki G (2004) *Direct Reasoning*, Ph. D. Dissertation, Clemson University.
- [36] Kulczycki G, Sitaraman M, Weide B W., Rountev A (2006) A specification-based approach to reasoning about pointers. In: *Proc ESEC/FSE SAVCBS '05 Workshop*. ACM Press.
- [37] Leavens G T, Baker A L, Ruby C (2006) Preliminary design of JML: a behavioral interface specification language for Java. *ACM Software Engineering Notes* 31:1-38.

- [38] Leavens G T, Abrial J, Batory D, Butler M, Coglio A, Fisler K, Hehner E, Jones C, Miller D, Peyton-Jones S, Sitaraman M, Smith D R, Stump A (2006) Roadmap for Enhanced Languages and Methods to Aid Verification. In: *Proc. GPCE 5*. ACM Press, New York, NY, pp 221-236.
- [39] Leinenbach D, Paul W, Petrova E (2005) Towards the formal verification of a C0 compiler: code generation and implementation correctness. *SEFM* 3:2-11.
- [40] Meyer B (1992) Applying design by contract. *Computer* 25:40-51.
- [41] Nipkow T, Paulson L C, Wenzel M (2002) Isabelle/HOL: A Proof Assistant for Higher-Order Logic. *LNCS 2283*. Springer-Verlag.
- [42] Owre S, Rajan S P, Rushby J M, Shankar N, Srivas M (1996) PVS: Combining specification, proof checking, and model checking. In: Alur R and Henzinger T A (eds.), *Proc. CAV*, LNCS 1102. Springer Verlag, pp 411-414.
- [43] Sitaraman M, Weide B W (1994) Component-based software using RESOLVE. *Software Engineering Notes* 19:21-67.
- [44] Sitaraman M, Weide B W, Ogden W F (1997) On the Practical Need for Abstraction Relations in Abstract Data Type Verification. *IEEE TSE* 23 (3), March 1997.
- [45] Sitaraman M, Atkinson S, Kulczycki G, Weide B, Long T, Bucci P, Pike S, Heym W, Hollingsworth J (2000) Reasoning about software-component behavior. *Proceedings of the 6th International Conference on Software Reuse*, LNCS 1844. Springer Berlin, pp 266-283.
- [46] Smith H, Roche K, Sitaraman M, Krone J, Ogden W F (2008) Integrating math units and proof checking for specification and verification. In: *Proc. SAVCBS*, pages 59-66.
- [47] Spec#. <http://research.microsoft.com/specsharp/> viewed 1 Dec 2008.
- [48] Weide B W, Sitaraman M, Harton H K, Adcock B, Bucci P, Bronish D, Heym W D, Kirschenbaum J, Frazier D (2008) Incremental benchmarks for software verification tools and techniques. In: *Proc. VSTTE 2008*, LNCS 5295. Springer-Verlag, pp 84-98.
- [49] Weide B W, Heym W (2001) Specification and Verification with References. In: *Proc. SAVCBS*.

- [50] Weide B W, Edwards S H, Harms D E, Lamb D A (1994) Design and Specification of Iterators Using the Swapping Paradigm. *IEEE TSE* 20(8):631-643.
- [51] Zee K, Kuncak V, and Rinard M C (2008) Full functional verification of linked data structures. In: *Proc. PLDI*. ACM, pp 349-361.