

# Getting Started with a Framework for Verifying Software Transactional Memory Algorithms

Mohsen Lesani   Victor Luchangco   Mark Moir

September 2012

## Abstract

This document is intended to help readers start to use our framework for formally verifying software transactional memory algorithms using the PVS verification system. A paper describing technical aspects of this framework, and some of the challenges encountered and lessons learned in using the framework, was presented at CONCUR 2012. This document describes more mundane aspects, such as how to explore the framework using PVS for readers not already familiar with it, and provides brief descriptions of the files contained in the framework.

## 1 Introduction

This document briefly describes a framework for formally verifying transactional memory (TM) algorithms that we presented at CONCUR 2012 [7], and gives some tips that we hope enables readers to start using this framework quickly and easily. Our framework is formalized using the PVS verification system [1, 11], and, for readers not already familiar with PVS, we describe how to get started using it quickly, and provide pointers to where to get it. Although we do not assume familiarity with PVS, this document is not a tutorial on PVS; such tutorials are available at the PVS website [1, 2].

## 2 Overview of the framework

The framework provides three kinds of files used by PVS:

- *specification files* (`.pvs` suffix) contain definitions and lemmas provided by the framework;
- *proof files* (`.prf` suffix) contain saved proofs of lemmas; and
- the `pvs-strategies` files defines *strategies* used in the proofs.

Most users will be interested primarily in the specification files. Proof files are generated and updated automatically by PVS, and should not be edited directly. Strategies are user-defined routines that capture common patterns in proofs. We do not discuss strategies further in this document; for more information, please refer to PVS documentation (e.g., [2]).

Definitions and lemmas are specified in the PVS language, which is a typed higher-order logic. PVS generates proof obligations for every lemma, as well as for *type-correctness conditions* (TCCs)





Start PVS in the directory containing the framework (i.e., the directory with the specification and proof files). For example, if this directory is `~/framework`, then go to `~/framework` and type `pvs` at the prompt:

```
~/framework$ pvs
```

This starts Emacs with PVS running (in the `*pvs*` buffer).

In the rest of this section, `C`, `M`, and `T` denote the control, meta, and tab keys respectively. PVS can be run interactively or in batch mode.

Although we discuss aspects of PVS in this section, it is neither a tutorial nor a definitive reference for PVS. Interested readers should consult the PVS documentation available on the PVS website [1], which includes a tutorial [2] and a guide to the prover [12]. However, note that this documentation is rather dated, and so is not entirely accurate for the latest version of PVS.

### 3.1 Running proofs

Open the file `Automata.pvs`, and look at lemma `lengthTruncate`. There is a saved proof for this lemma. (Proofs of lemmas in `Automaton.pvs` are saved in `Automaton.prf`.) To see this proof, place the cursor on `lengthTruncate` lemma and type

```
M-x show-proof
```

A buffer named `Proof` (henceforth, the *proof buffer*) is opened, displaying the saved proof. (A proof in PVS is a script consisting of a sequence of commands to the prover.)

The PVS prover runs all proofs in the `*pvs*` buffer. To run the saved proof of a lemma, place the cursor on the lemma (in the specification file) and press

```
C-c p
```

At the prompt, type `yes` when you are asked whether to rerun the existing proof. (Type `no` to write a new proof.) This starts the proof in the `*pvs*` buffer. As the proof is correct, it finishes successfully. If there were no saved proofs for the lemma, this step would be skipped and you would be immediately prompted to prove the lemma in the `*pvs*` buffer.

We can also step through the proof: Go to the lemma `lengthTruncate` again and press

```
C-c C-p s
```

This loads the proof buffer with the saved proof, which consists of four commands. It also makes the `*pvs*` buffer current, and prompts for a prover command. Above the prompt, the current sequent is displayed. We must establish that the conjunction of formulas above the turnstile implies the disjunction of the formulas below it.

To execute the next command from the proof buffer, press

```
T 1
```

This executes the command following the cursor in the proof buffer, and advances the cursor in the proof buffer to the next command. Thus, it can be used repeatedly to run consecutive commands. To run the next  $k$  command consecutively, press

```
C-u k T 1
```

The proof buffer can be edited, and the cursor can be moved inside the proof buffer to mark a different command as the next command.

Note that `T 1` executes whatever command follows the cursor in the proof buffer. Thus, one can load the proof of one lemma into the proof buffer (e.g., by `M-x show-proof`), and then use it to prove a different lemma, provided that the proof buffer is not overwritten (e.g., by pressing `C-c C-p s` to start proving the other lemma).

To issue a command directly (rather than running one from the proof buffer), either type it at the prompt in the `*pvs*` buffer or use a command shortcut. Shortcuts for some of the commands we use most frequently are:

<code>T a</code> (assert)	<code>T g</code> (ground)	<code>T p</code> (prop)	<code>T =</code> (decompose-equality)
<code>T e</code> (expand)	<code>T G</code> (grind)	<code>T s</code> (split)	<code>T ?</code> (inst?)
<code>T f</code> (flatten)	<code>T l</code> (lift-if)	<code>T S</code> (skosimp)	

See the PVS prover guide [12] for a description of these commands.

Some commands generate multiple proof obligations as subgoals, leading proofs to have a tree structure. The `(postpone)` command (shortcut: `T P`) causes the prover to switch to the next incomplete branch in the proof tree. The `(undo)` command (shortcut: `T u`) removes the last command in the current branch of the proof tree. Note that this is *not* the previous command issued if that command completed some branch of the tree. If the undone command had some completed branches, those branches are also removed from the tree. A mistaken `(undo)` command can be reversed using `(undo undo)`, but only if this is done immediately afterwards (i.e., with no intervening command).

PVS provides a graphic display of the proof tree. To see the current proof tree, type

`M-x x-show-current-proof`

This tree is updated as you issue more commands, so that it always reflects the current state of the proof. This is helpful to see proof branches in large proofs. Clicking on the turnstiles that appear between commands opens a window that displays the sequent at that point in the proof. Long commands are abbreviated in the proof tree; clicking on them opens a window showing the full command.

Certain commands, particularly `(grind)`, may run for a long time—indeed, they may never terminate (or at least run longer than we have been willing to wait). You can interrupt such a command and return to the point immediately before executing it by typing

`C-c C-c C-d`

To quit a proof, use the `(quit)` command (shortcut: `T q`). If you have issued a sequence of commands other than the saved proof, you are prompted whether the proof should be saved. In this case, type `no` to preserve the existing proof.

A theory (and any theories it imports) must be type-checked before it can be proved. This is done automatically when a proof in the theory is attempted, but it can also be done explicitly by typing

`M-x typecheck-prove`

TCCs are automatically generated constraints (lemmas) to enforce type soundness. Some of them are automatically discharged and some are not. To see the TCCs of this theory, type

`M-x show-tccs`

or

`C-c C-q s`

### 3.2 Rerun the proofs in batch

To run all the saved proofs for a theory, place the cursor anywhere within the theory and press

```
C-c C-p t
```

After executing all the proofs, the status of all lemmas and TCCs of the theory is displayed in a new buffer called **PVS Status**. The status of a lemma or TCC may be “untried”, “unfinished”, “proved - incomplete” or “proved - complete”. A lemma is “proved - incomplete” if it was proved but used lemmas that are not “proved - complete”. To get this status again later, type

```
M-x status-proof-theory
```

To prove all the theories that a theory directly or indirectly imports, type

```
C-c C-p i
```

which runs all the proofs and then brings a summary of the status of the theories in the import chain. To prove all theories in the entire framework, invoke this command on the **Main** theory. To see this status again later, type

```
M-x status-proofchain
```

### 3.3 Get information

To quickly find the definition of a function in a specification file, place the cursor on a usage of this function and press

```
M-.
```

For example, the `lengthTruncate` lemma uses the `truncate` function. In this case, the definition is right above the lemma but it could be far away, even in a different (imported) theory. Placing the cursor on the call to `truncate` in `lengthTruncate` and pressing `M-.` opens a new buffer with the definition of `truncate`. Conversely, to see where a function is used, place the cursor over the function and press

```
M-;
```

Placing the cursor on the `truncate` function and pressing `M-;` opens a new buffer with a list of the places where `truncate` is used, including the lemma `lengthTruncate`.

To see the theories in the PVS prelude, type

```
M-x view-prelude-file
```

To get the list of all PVS prover commands, type

```
M-x help-pvs-prover
```

In PVS, definitions can be *overloaded*; that is, there can be many definitions for the same name. PVS resolves overloading by context, and causes type-checking to fail if it cannot do so. When displaying sequents, however, it typically just shows the plain overloaded names without the context. Thus, two expressions that appear to be identical might differ because an overloaded name might be resolved differently. To see how user-defined overloading is resolved, type

```
M-x show-expanded-sequent
```

This does not show the resolution of names overloaded in the prelude. To see these, press

```
C-u M-x show-expanded-sequent
```

(Note that infix operators are expanded into binary functions, so that an expanded sequent might look quite different from the unexpanded sequent.)

### 3.4 Issues with PVS

We encountered some difficulties in using PVS, many of which were resolved with help from Sam Owre. For some others, we developed workarounds, which we discuss here.

First, sometimes PVS generates extra TCCs (for which it cannot, of course, find any proofs) or complains about names already being defined. We believe that this has to do with how PVS stores intermediate computations, both internally and in `.bin` files. When this happens, we find it easiest to “reboot” PVS as follows: Kill the `*pvs*` buffer by typing

```
C-x k *pvs*
```

(Emacs will ask whether you are sure you want to kill the PCS process.) Then delete any `.bin` files in the directory. Then restart PVS by typing

```
M-x pvs
```

We notice that PVS typically gets confused about names in theories that use theory aliases, and intend to excise such use from our files. However, our current version of the framework still has many such uses.

Second, we discovered a case in which PVS failed to generate certain TCCs, allowing unsound proofs. Fortunately, Sam Owre fixed this problem and provided us a patch (see Appendix A.1), which should be included in the next release of PVS.

Third, PVS allows theories to be organized into “libraries” (i.e., directories), to provide an extra level of modularity. We would like to exploit this functionality by having separate libraries for the basic automata theory, the TM specifications, and the automata related to NOrec, and theories that extend the PVS prelude theories of finite sequences and partial functions. (This classification is embodied in the different “main” files we provide.) However, we have not been able to get PVS to consistently confirm that all proofs are complete when they depend on lemmas defined in theories imported from other libraries. Thus, for now, we keep all the files in a single library.

## A Installing PVS

PVS is developed and maintained by SRI International. It can be downloaded from their website [1], along with instructions on how to install it. Our framework was (and is being) developed using PVS 5.0 running on Linux on a 32-bit Intel machine. For licensing reasons, we use the SBCL version. To install this version, download `pvs-5.0-ix86-Linux-sbclisp` from the PVS website, and unzip and untar it in the directory in which you want to install PVS. Let’s say that directory is `~/PVS`:

```
~/PVS$ tar -zxvf pvs-5.0-ix86-Linux-sbclisp.tgz
```

Now, run the relocation script

```
./~/PVS/bin/relocate
```

and add the PVS folder to your path

```
~/PVS$ export PATH=$PATH:~/PVS
```

You should now be able to run pvs

```
~/PVS$ pvs
```

## A.1 PVS patch

As mentioned above, there is a bug in PVS that caused it to fail to generate certain type-correctness conditions. Sam Owre provided the following patch, which should be added to the `.emacs` file.

```
(in-package :pvs)

(defmethod make-implicit-conversion ((conv expr) ctype ex)
  (let ((nexpr (copy-untyped ex)))
    (change-class ex 'implicit-conversion)
    (setf (argument ex) nexpr)
    (setf (types nexpr) nil)
    (setf (operator ex)
          (raise-actuals (copy (expr (from-conversion ctype))) 1))
    (typecheck* ex nil nil nil)))
```

## A.2 Shortcuts

Some users, especially those not familiar with Emacs, may find the Emacs commands confusing. Such users may find it useful to remap many common commands of Emacs and PVS to more familiar keys. Such a remapping is done by the `.emacs` file that can be retrieved from

<http://cs.ucla.edu/~lesani/downloads/.emacs>

Follow the simple instructions in the `.emacs` file to install it. The shortcut keys it provides are listed at the beginning of the file `.emacs` file.

## References

- [1] PVS website. <http://pvs.csl.sri.com/>.
- [2] Judy Crow, Sam Owre, John Rushby, Natarajan Shankar, and Mandayam Srivas. A tutorial introduction to PVS. In *Workshop on Industrial-Strength Formal Specification Techniques*, April 1995.
- [3] Luke Dalessandro, Michael Spear, and Michael Scott. NOrec: Streamlining STM by abolishing ownership records. In *Symposium on Principles and Practice of Parallel Programming*, January 2010.
- [4] Simon Doherty, Lindsay Groves, Victor Luchangco, and Mark Moir. Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing*, 2012.
- [5] Rachid Guerraoui and Michal Kapalka. On the correctness of transactional memory. In *Symposium on Principles and Practice of Parallel Programming*, February 2008.
- [6] Rachid Guerraoui and Michal Kapalka. *Principles of Transactional Memory*. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool, 2010.
- [7] Mohsen Lesani, Victor Luchangco, and Mark Moir. A framework for formally verifying software transactional memory algorithms. In *International Conference on Concurrency Theory (CONCUR)*, September 2012.



- [8] Mohsen Lesani, Victor Luchangco, and Mark Moir. Putting opacity in its place. In *Workshop on Theory of Transactional Memory*, July 2012.
- [9] Nancy Lynch and Mark Tuttle. Hierarchical correctness proofs for distributed algorithms. In *Symposium on Principles of Distributed Computing*, pages 137–151, August 1987.
- [10] Nancy Lynch and Frits Vaandrager. Forward and backward simulations part I: Untimed systems. *Information and Computation*, 121(2):214–233, September 1995.
- [11] Sam Owre, Natarajan Shankar, and John Rushby. PVS: A prototype verification system. In *International Conference on Automated Deduction*, pages 748–752, June 1992.
- [12] N. Shankar, S. Owre, J. M. Rushby, and D. W. J. Stringer-Calvert. *PVS Prover Guide*, November 2001.