

# Automating Theorem Proving with SMT

K. Rustan M. Leino

Microsoft Research, Redmond, WA, USA  
leino@microsoft.com

**Abstract.** The power and automation offered by modern satisfiability-modulo-theories (SMT) solvers is changing the landscape for mechanized formal theorem proving. For instance, the SMT-based program verifier Dafny supports a number of proof features traditionally found only in interactive proof assistants, like inductive, co-inductive, and declarative proofs. To show that proof tools rooted in SMT are growing up, this paper presents, using Dafny, a series of examples that illustrate how theorems are expressed and proved. Since the SMT solver takes care of many formal trivialities automatically, users can focus more of their time on the creative ingredients of proofs.

## 1 Introduction

A growing number of theorems about mathematics, logic, programming-language semantics, and computer programs are formalized and proved using mechanized proof assistants. Examples of such proof assistants are ACL2 [24], Agda [9,35], Coq [6], Guru [40], HOL Light [19], Isabelle/HOL [34], PVS [36], and Twelf [37]. The assistants vary in their level of expressivity and automation as well as in the size of their trusted computing base. Satisfiability-modulo-theories (SMT) solvers (for example, Alt-Ergo [7], CVC3 [3], OpenSMT [10], Simplify [14], and Z3 [13]) are collections of (semi-)decision procedures for certain theories. SMT solvers provide a high degree of automation and have, in the last couple of decades, undergone impressive improvements in power. Therefore, it has become increasingly common for proof assistants to use SMT solvers as subroutines, as is done for example in PVS [36] and in Isabelle/HOL's Sledgehammer tactic [8].

Although general proof assistants can be used to verify the correctness of computer programs, there are also some verification tools dedicated to verifying programs. These include Chalice [30], Dafny [26], F\* [41], Frama-C [12], Hi-Lite Ada [18], KeY [4], KIV [39], Pangolin [38], Spec# [2], VCC [11], VeriFast [21], and Why3 [17]. Many of these use as their underlying reasoning engine an SMT solver, typically accessed via an intermediate verification language like Boogie [1] or Why [16]. This tool architecture facilitates automation, and it also tends to move the user's interaction with the tool from the formula level (like in general proof assistants) to the program level. This lets users express necessary proof ingredients in program-centric declarations like preconditions or loop invariants. We might therefore refer to this kind of program verifier as *auto-active*—a mix of automatic decision procedures and user interaction at the program level [28].

Just as some proof assistants have incorporated special tactics to better handle program verification (e.g., Ynot [33]), auto-active program verifiers are incorporating features to better support user-guided mathematical proofs (e.g., VeriFast [22] and Dafny [26]). While such program verifiers do not yet achieve the full expressivity of some proof assistants and generally have a much larger trusted computing base, their automation can be remarkable. This automation matters, since it affects the amount of human time required to use the tool. Curiously, these SMT-based tools are primarily program verifiers, but that seems to have happened more by serendipity; one can easily imagine similar SMT-based tools that focus on mathematics rather than on programs.

In this paper, and in the invited talk that the paper accompanies, I argue and showcase that the future may hold room for proof assistants that are entirely based on SMT solving. I focus on the programming language and auto-active program verifier Dafny [26], which supports proof features like induction [27], co-induction [29], and declarative calculations [31]. The paper is a collection of examples that give an idea of what these proof features can do and how they are represented in the Dafny input.

Sec. 2 defines a type and a function that will be used throughout the examples. Sec. 3 states a lemma and proves it by induction. It also shows a proof calculation. Sec. 4 then turns to the infinite by considering co-inductive declarations and proofs. Sec. 5 combines the inductive and co-inductive features into a famous “filter” function, and Sec. 6 proves a theorem about filters by simultaneously applying induction and co-induction. Sec. 7 summarizes the examples and concludes.

## 2 A Types and a Function

As a basis for all the examples in this paper, let us define a type `Stream` whose values are infinite lists. Such a type is called a *co-inductive datatype*, but this fancy name need not cause any alarms.

```
codatatype Stream<T> = Cons(head: T, tail: Stream)
```

The stream type is parameterized by the type of its elements, `T`. The stream type has one constructor, `Cons`. The names of the parameters to the constructor (`head` and `tail`) declare destructors. For example, for any stream `s`, we have

```
s = Cons(s.head, s.tail)
```

The type of `tail` is `Stream<T>`, but here and in other signatures, the type argument `T` can be supplied automatically by Dafny, so I omit it.

Next, let us declare a function that returns a suffix of a given stream. In particular, `Tail(s, n)` returns `s.tailn`. Here is its inductive definition:

```
function Tail(s: Stream, n: nat): Stream
{
  if n = 0 then s else Tail(s.tail, n-1)
}
```

Each function is checked for well-definedness. For a recursive function, this includes a check of well-foundedness among the recursive calls. Well-foundedness is checked

using a *variant function*, which conceptually is evaluated on entry to a call. If the value of the variant function is smaller for the callee than for the caller, well-foundedness follows. If no variant function is explicitly supplied, Dafny guesses one. The guess is the lexicographic ordering on the tuple of the function’s arguments, omitting arguments whose types have no ordering, like co-inductive datatypes. For `Tail`, Dafny (correctly) guesses the variant function `n` and, using it, checks that `Tail`’s recursion is indeed well-founded. No further input is required from the user—the tool automatically initiates the check—and since the check succeeds, the user is not bothered by any messages from the tool (except for a small indication in the margin of the integrated development environment that for a split second shows that the verifier is active).

### 3 An Inductive Proof

In order to show how lemmas and proofs are set up, let us consider an alternative definition of `Tail`:

```
function Tail_Alt(s: Stream, n: nat): Stream
{
  if n = 0 then s else Tail_Alt(s, n-1).tail
}
```

and let us prove that `Tail` and `Tail_Alt` give the same result.

A lemma is expressed as a *method* in the programming language, that is, as a code procedure with a pre- and postcondition. As usual (e.g., [20]), such a method says that for any values of the method parameters that satisfy the method’s precondition, the method will terminate in a state satisfying the postcondition.<sup>1</sup> This corresponds to what is done in mathematics for a lemma: a lemma says that for any values of the lemma parameters that satisfy the lemma’s antecedent, the lemma’s conclusion holds. So, we express our lemma about `Tail` and `Tail_Alt` by declaring the following method:

```
ghost method Tail_Lemma(s: Stream, n: int)
  requires 0 ≤ n;
  ensures Tail(s, n) = Tail_Alt(s, n);
```

The precondition of this method (keyword **requires**) says that `n` is a natural number (which, alternatively, we could have indicated by declaring the type of `n` to be **nat**). The postcondition (keyword **ensures**) gives the property we want to prove. The designation of the method as a *ghost* says that we do not want the Dafny compiler to emit any executable code. In other words, a ghost method is for the verifier only; the compiler ignores it.

To get the program verifier to prove the lemma, we supply a method body and let the verifier convince itself that all code control paths terminate and establish the postcondition. The body we supply typically consists of **if** statements and (possibly recursive) method calls, but other statements (e.g., **while** loops) can also be used. A recursive call

---

<sup>1</sup> In general, methods can have effects on the memory of the program. Such effects are declared with a **modifies** clause in the method specification. However, since there is no need for such effects here, I ignore further discussion of them.

corresponds to invoking an inductive hypothesis, because the effect on the proof is to obtain the proof goal (stated in the postcondition) for the callee’s arguments, which in a well-foundedness check are verified to be “smaller” than the caller’s arguments. It is common to supply assertions that act as in-place lemmas: a statement **assert**  $Q$  tells the verifier to check that the boolean condition  $Q$  holds, after which the verifier can make use of that condition. A more systematic way to direct the verifier is provided by Dafny’s **calc** statement, whose verified calculations take the form of human-readable equational proofs [31].

Here is a proof of the lemma:

```
ghost method Tail_Lemma(s: Stream, n: int)
  requires 0 ≤ n;
  ensures Tail(s, n) = Tail_Alt(s, n);
{
  if n < 2 {
    // def. of Tail and Tail_Alt
  } else {
    calc {
      Tail(s, n);
    = // def. Tail, since n ≠ 0
      Tail(s.tail, n-1);
    = { Tail_Lemma(s.tail, n-1); } // induction hypothesis
      Tail_Alt(s.tail, n-1);
    = // def. Tail_Alt, since n-1 ≠ 0
      Tail_Alt(s.tail, n-2).tail;
    = { Tail_Lemma(s.tail, n-2); } // induction hypothesis
      Tail(s.tail, n-2).tail;
    = // def. Tail, since n-1 ≠ 0
      Tail(s, n-1).tail;
    = { Tail_Lemma(s, n-1); } // induction hypothesis
      Tail_Alt(s, n-1).tail;
    = // def. Tail_Alt, since n ≠ 0
      Tail_Alt(s, n);
    }
  }
}
```

The method body provides two code paths. For the  $n < 2$  branch, the verifier can prove the postcondition by unwinding the definitions of each of `Tail` and `Tail_Alt` once or twice (which the verifier is willing to do automatically). The `else` branch uses a **calc** statement with a number of equality-preserving steps, each of which is verified. Some steps are simple and need no further justification; the code comments give explanations for human consumption. Other steps are justified by *hints*, which are given as code blocks (in curly braces). Here, each hint makes a recursive call to `Tail_Lemma`, which in effect invokes the induction hypothesis.

In more detail, for each step in a calculation, the verifier checks that the equality entailed by the step is provable after the code in the associated hint (if any). In the

calculation above, each of the provided hints consists of a single recursive call. As usual in program verification, the verifier thus checks that the precondition of the callee is met, checks that the variant function is decreased for the call (to ensure that the recursion will terminate), and can then assume the postcondition of the callee. For example, the postcondition that holds after the first recursive call is:

$$\text{Tail}(s.\text{tail}, n-1) = \text{Tail\_Alt}(s.\text{tail}, n-1)$$

which is essentially the induction hypothesis for  $s, n := s.\text{tail}, n-1$ . Since no variant function is supplied explicitly, Dafny guesses  $n$ , which it verifies to decrease. Thus, the recursion—and indeed, the induction—is well-founded.

For brevity, the equality signs between the lines in the calculation can be omitted. Or, if desired, they can be replaced by different operators, like  $\implies$ ,  $\Leftarrow$ , or  $\leftarrow$ .

The calculation in the example gives more detail than the Dafny verifier needs, but, as given, yields a presentation of the proof that is better suited for a human. In fact, the proof calculation is quite readable; it looks almost identical to how one would write an equational-style proof by hand. For a comparison with other styles of declarative proofs, like Isar [42], and with tactic-based proofs, see [31].

## 4 Co-recursion and a Co-inductive Proof

In this section, we consider how values of a co-datatype are constructed and how one states and proves properties of such values.

Values of co-inductive datatypes may be of an infinite nature. For example, a stream represents an infinite list of elements. Here is a function that defines such a value, namely the stream whose elements are the integers from  $n$  upward in increasing order:

```
function Up(n: int): Stream<int>
{
  Cons(n, Up(n+1))
}
```

It may look as if invocations of `Up` will never terminate, but the self-call of `Up` is identified by Dafny as being *co-recursive*, because it is positioned as an argument to a co-datatype constructor. Co-recursive calls are compiled into lazily evaluated code, so that the arguments to the constructor are not evaluated until their values are used by the executing program (if ever). Consequently, for a co-recursive call, there is no need for the verifier to enforce a decrease of a variant function.

Here is another function on streams:

```
function Prune(s: Stream): Stream
{
  Cons(s.head, Prune(s.tail.tail))
}
```

It defines a stream consisting of half of the elements of the given stream: every other element, starting with the first. Note that the self-call to `Prune` is co-recursive, so the verifier does not need to check termination.

To define a property of a co-inductive datatype, one uses a *co-predicate*. For example, the following co-predicate holds for streams that consist of even integers:

```
copredicate AllEven(s: Stream<int>)
{
  s.head % 2 = 0 ∧ AllEven(s.tail)
}
```

Co-predicates are defined by greatest fix-points, that is, as the greatest solutions of the recursive equations to which their definitions give rise. Applied to the example, this means that `AllEven(s)` evaluates to **true** as long as there is no suffix `t` of `s` such that `t.head % 2 ≠ 0`. Eager evaluation of a co-predicate may fail to terminate and lazy evaluation would not be meaningful, so co-predicates are always ghost. In other words, they are never part of executing code, but they can be used to describe and reason about executing code.

We have now seen three features from the quartet of co-inductive features in Dafny: co-datatypes define possibly infinite data structures, co-recursive function calls make it possible to define values of co-datatypes, and co-predicates define properties of co-datatypes. The fourth feature is *co-methods*, whose purpose is to enable co-inductive proofs. Let us consider an example.

We will state a theorem that for any even `n`, `Prune(Up(n))` consists only of even integers. Because we intend to prove the theorem by co-induction, we use a co-method:

```
comethod Theorem(n: int)
  requires n % 2 = 0;
  ensures AllEven(Prune(Up(n)));
{
  Theorem(n+2);
}
```

Ignoring the issue of termination, this proof can be understood in the same manner as inductive proofs: `AllEven` says something about the head of the stream `Prune(Up(n))`, which is proved automatically. It also says something about the tail of the stream, which follows from the postcondition of the call `Theorem(n+2)` and the definitions of the functions involved. To make this argument more explicit, the call could have been preceded by the following calculation (where, for brevity and variety, I have chosen to omit the optional equality signs between lines in the left margin):

```
calc {
  Prune(Up(n)).tail;
  Prune(Up(n).tail.tail);
  { assert Up(n).tail.tail = Up(n+2); }
  Prune(Up(n+2));
}
```

In contrast to methods, whose recursive calls are checked to terminate (by checking that they decrease the variant function), calls to co-methods are always allowed. In other words, the co-induction hypothesis can always be obtained; however, the *use* of it is restricted. Intuitively, the co-induction hypothesis can be used to discharge only

those conjuncts that show up after one unwinding of the co-predicate in the co-method's postcondition. I will give some details about this in Sec. 6.

For example, suppose the body of co-method `Theorem` were replaced by the call `Theorem(n)`. With unrestricted use of the postcondition of this call, the co-induction hypothesis obtained would trivially prove the theorem itself. However, because the co-induction hypothesis can be used only on conjuncts from an unwinding of the postcondition, the call `Theorem(n)` provides no benefit here.

## 5 A Filter Function

Let us now consider a more difficult function definition, namely that of a *filter* function on streams. For any stream `s`, we want the filter function to return the stream consisting of those elements of `s` that satisfy some predicate `P`. A filter function like this is used, for example, in the prime number sieve of Eratosthenes (*cf.* [5,25,15]).

Conceptually, the filter function and all related lemmas are parameterized by the predicate `P`. Lacking the higher-order features necessary to take `P` as a parameter, we represent an arbitrary predicate by declaring a (here, global) generic predicate without a defining body:<sup>2</sup>

```
predicate P<T>(x: T)
```

The definition of `Filter` has the following form:

```
function Filter(s: Stream): Stream
  //... specification to be written...
{
  if P(s.head) then
    Cons(s.head, Filter(s.tail))
  else
    Filter(s.tail)
}
```

The first branch of this definition is fine, because its call to `Filter` is co-recursive. However, the other call to `Filter` is not co-recursive, so it is subject to a termination check. This makes sense, because if the given stream has no elements that satisfy `P`, then `Filter` would never terminate in its computation to produce the next element of the resulting stream. Note, thus, how Dafny allows one function to be involved in both recursive and co-recursive calls. Next, we will consider how to deal with the termination of the recursive call.

To avoid non-termination, we must restrict `Filter`'s input to streams that contain infinitely many elements that satisfy `P`. We give the following definitions:

```
predicate HasAnother(s: Stream)
{
```

---

<sup>2</sup> Dafny allows such a body-less predicate to be placed in a *module*. Other modules can then be declared as *refinements* of this module, and each refinement module can give its own specific definition of the predicate.

```

  ∃ n • 0 ≤ n ∧ P(Tail(s, n).head)
}
copredicate AlwaysAnother(s: Stream)
{
  HasAnother(s) ∧ AlwaysAnother(s.tail)
}

```

Predicate `HasAnother(s)` says that, after some finite prefix of `s`, there is an element that satisfies `P`, and `AlwaysAnother(s)` says that `HasAnother` holds at every point in the stream. We can now restrict the input to `Filter` by adding a precondition:

```
requires AlwaysAnother(s);
```

Even with this precondition, the verifier complains that it cannot prove termination. To remedy the situation, we supply a variant function explicitly. As the variant function, we will use the length of the non-`P` prefix of `s`, that is, the number of steps to the next element satisfying `P`. Using `StepsToNext(s)` to denote that number of steps, we add to the specification of `Filter` the following clause:

```
decreases StepsToNext(s);
```

Given a stream that satisfies `AlwaysAnother`, function `StepsToNext` returns a natural number. It is tempting to define it with a body like

```
if P(s.head) then 0 else 1 + StepsToNext(s.tail)
```

but to prove that this recursive call to `StepsToNext` terminates, we would need a variant function like `StepsToNext` itself. Instead, we find a number of steps that will yield some `P` element, and then we use this number as an upper bound in a linear search to the first `P` element:

```

function StepsToNext(s: Stream): nat
  requires AlwaysAnother(s);
{
  var n :| 0 ≤ n ∧ P(Tail(s, n).head);
  Steps(s, n)
}
function Steps(s: Stream, n: nat): nat
  requires P(Tail(s, n).head);
  ensures P(Tail(s, Steps(s, n)).head);
  ensures ∀ i • 0 ≤ i < Steps(s, n) ⇒ ¬P(Tail(s, i).head);
{
  if P(s.head) then 0 else 1 + Steps(s.tail, n-1)
}

```

These definitions require some explanation.

The “let such that” expression `var x :| Q; E` evaluates to `E` in which all free occurrences of `x` are bound to a value that satisfies `Q`. The expression is well-defined only if there exists a value for `x` that satisfies `Q`. In `StepsToNext`, this proviso follows from the precondition `AlwaysAnother(s)`. Note that `n` may be set to any number of steps that will reach a `P` element in `s`, not necessarily the smallest.



The specification of the auxiliary function `Steps` requires `s` to reach a `P` element in `n` steps. It ensures that the result value, which in the **ensures** clause is denoted by `Steps(s, n)`, is not only a number of steps that reaches a `P` element (first **ensures** clause) but also the smallest such number (second **ensures** clause).

The body of `Steps` encodes a straightforward linear search.

To prove the recursive call in the body of `Filter(s)` to be well-founded, the verifier checks that the given variant function decreases, that is,

$$\text{StepsToNext}(s.\text{tail}) < \text{StepsToNext}(s)$$

This condition rests on the fact that `StepsToNext` returns the smallest number of steps to reach a `P` element, which is spelled out by the postcondition of `Steps`. Note that `StepsToNext` does not need to declare such a postcondition, because Dafny unwinds the definition of `StepsToNext` and obtains an expression in terms of `Steps`. Since `Steps` is recursive, the needed property is not evident from any bounded number of unwindings, so the presence of the postcondition essentially facilitates an inductive argument.

Finally, rather than introducing the auxiliary function `Steps`, one could consider replacing the body of `StepsToNext` with one whose let-such-that condition is stronger:

```
var n :| 0 ≤ n ∧ P(Tail(s, n).head) ∧
        ∀ i • 0 ≤ i < n ⇒ ¬P(Tail(s, i).head);
n
```

However, Dafny is unable to prove the existence of such an `n` directly from the precondition `AlwaysAnother(s)`. The use of `Steps` is one way to set up the necessary inductive argument.

## 6 A Property of Filter

The interesting property to prove about `Filter(s)` is that it returns the subsequence of `s` that consists of exactly those elements that satisfy `P`. The notion of such a subsequence can be divided up into the property that `Filter` returns the right set of elements:

$$\forall x \bullet x \in \text{Filter}(s) \iff x \in s \wedge P(x)$$

(where I have taken the liberty of using operator  $\in$  as if stream were sets) and the property that `Filter(s)` preserves the order of elements in `s`. Let us look at one possible way to state and prove the latter.

To simplify matters, let us suppose that there is a function `Ord` from the elements of streams to the integers.

```
function Ord<T>(x: T): int
```

Using a co-predicate, we define what it means for a stream's elements to be strictly increasing:

```
copredicate Increasing(s: Stream)
{
  Ord(s.head) < Ord(s.tail.head) ∧ Increasing(s.tail)
}
```

Now we can state the order-preservation theorem that we want to prove:

```
ghost method Theorem_FilterPreservesOrdering(s: Stream)
  requires AlwaysAnother(s)  $\wedge$  Increasing(s);
  ensures Increasing(Filter(s));
```

This theorem is not the most general order-preservation theorem we can state, but it suffices for our purpose of showing an interesting proof.

To prove the theorem, we introduce an alternative definition of Increasing:

```
copredicate IncrFrom(s: Stream, low: int)
{
  low  $\leq$  Ord(s.head)  $\wedge$  IncrFrom(s.tail, Ord(s.head) + 1)
}
```

The two definitions are interchangeable, as the following two lemmas show.

```
comethod Lemma_Incr0(s: Stream, low: int)
  requires IncrFrom(s, low);
  ensures Increasing(s);
{
}
comethod Lemma_Incr1(s: Stream)
  requires Increasing(s);
  ensures IncrFrom(s, Ord(s.head));
{
  Lemma_Incr1(s.tail);
}
```

The co-inductive proof of Lemma\_Incr0 is done automatically, whereas the other requires an explicit appeal to the co-induction hypothesis.

We can now write the theorem in terms of IncrFrom:

```
comethod Lemma_FilterPreservesIncrFrom(s: Stream, low: int)
  requires AlwaysAnother(s)  $\wedge$  IncrFrom(s, low)  $\wedge$  low  $\leq$  Ord(s.head);
  ensures IncrFrom(Filter(s), low);
  decreases StepsToNext(s);
{
  if P(s.head) {
    Lemma_FilterPreservesIncrFrom(s.tail, Ord(s.head) + 1);
  } else {
    Lemma_FilterPreservesIncrFrom#[_k](s.tail, low);
  }
}
```

The proof of this lemma is interesting because it uses co-induction and induction together. The first branch of the **if** statement makes an appeal to the co-induction hypothesis. Dafny will actually fill it in automatically, so the proof also goes through with that call omitted. In the else branch, we cannot use the co-induction hypothesis, because, as discussed above, the co-induction hypothesis can be used only after one unwinding

of the proof goal. To make use of the lemma's postcondition for `s.tail` directly, we instead make a recursive call to the lemma. Syntactically, this is achieved by the characters “#[\_k]”. The recursive call gives rise to a proof obligation of termination, which is addressed by the explicit **decreases** clause.

Co-methods are used to establish the validity of co-predicates (including equality on co-datatype values, which is a built-in co-predicate). These co-inductive proof obligations are actually carried out by induction, in conjunction with a meta-theorem. For any co-predicate  $Q(x)$ , let the *prefix predicate*  $Q\#[_k](x)$  denote the first  $_k$  unrollings of  $Q$ , defined inductively. For example, the prefix predicate for co-predicate `AllEven` is:<sup>3</sup>

```
predicate AllEven#[_k: nat](s: Stream<int>)
{
  if _k = 0 then
    true
  else
    s.head % 2 = 0 ^ AllEven#[_k-1](s.tail)
}
```

Similarly, for each co-method  $M(x)$ , Dafny generates a *prefix method*  $M\#[_k](x)$ , where each call  $M(E)$  in the co-method's body is turned into a call  $M\#[_k-1](E)$  in the corresponding prefix method. In more detail, the following co-method:

```
comethod M(x: T)
  ensures Q(x);
  decreases D(x);
{
  ... M(E); ...
}
```

is turned into:

```
ghost method M#[_k: nat](x: T)
  ensures Q#[_k](x);
  decreases _k, D(x);
{
  if _k ≠ 0 {
    ... M#[_k-1](E); ...
  }
}
```

Any explicit prefix-method call in the body of  $M$  (like the one in the else branch of the filter lemma co-method above) is left unchanged in the corresponding prefix method. A recursive call to  $M\#[K]$  where  $K < _k$  corresponds to obtaining the co-induction hypothesis (for use after  $_k - K$  unwindings of the co-predicate in the proof goal), whereas a call to  $M\#[_k]$  is just an ordinary recursive call corresponding to the induction

---

<sup>3</sup> Prefix predicates are declared automatically. The made-up declaration syntax shown here is suggestive of how prefix predicates are actually invoked, with the unrolling-depth argument in square brackets, set apart from the other arguments.

hypothesis. By verifying the inductive prefix method for any  $\_k$ , the postcondition of the co-method follows on account of the following meta-theorem [32,29]:

$$\forall x: T \bullet Q(x) \iff \forall \_k: \mathbf{nat} \bullet Q\#[\_k](x)$$

For more details, see [29].

Finally, the proof of `Theorem_FilterPreservesOrdering` is given as follows:

```
{
  Lemma_Incr1(s);
  Lemma_FilterPreservesIncrFrom(s, Ord(s.head));
  Lemma_Incr0(Filter(s), Ord(s.head));
}
```

## 7 Conclusion

In this paper, I have conveyed a flavor of Dafny’s proof features by showing examples of inductive and co-inductive definitions, proofs by induction and by co-induction, as well as human-readable proofs. These are features that until recently were confined to interactive proof assistants, but they can now be supported by auto-active verifiers.

To try the examples in the Dafny tool,<sup>4</sup> the only input given to the tool are the lines shown in this paper—no additional proof tactics need to be supplied.

The given examples showcase the high degree of automation that is possible in a tool powered by an SMT solver and designed to keep the interaction at the problem level (and not, for example, at the level of defining and using necessary prover tactics). Users are not bothered with trivial details (like the associativity of logical and arithmetic operators) and the human involvement to prove that user-defined functions are mathematically consistent is small. Even the tricky recursive call of `Filter` is solved by defining and using `StepsToNext` as a variant function, which does not require an excessive amount of human effort. When more information is needed, human-readable calculations can be used, putting proofs in a format akin to what may be done by hand.

For each function and method in the examples shown, the verifier needs to spend only a small fraction of a second. This makes performance good enough to be running the verifier continuously in the background of the integrated development environment, which is what Dafny does. Most changes of the program text yield a near-instant response, which is important when developing proofs. Note that performance is at least as important for failed proofs as for successful proofs, because failed proofs happen on the user’s time (see [28] for some research directions for auto-active verification environments).

In the future, I expect a higher degree of automation to become available in proof assistants. For tools like Dafny that already provide a high degree of automation, I expect to see a richer set of features (for example, higher-order functions, drawing inspiration from Pangolin [38], Who [23], and F\* [41], and user-defined theories, drawing inspiration from Coq [6] and Why3 [17]) as well as work that will seek to reduce the currently large trusted computing base for SMT-based verifiers.

<sup>4</sup> Dafny can be installed from <http://dafny.codeplex.com>. It can also be run directly in a web browser at <http://rise4fun.com/dafny>

**Acknowledgments.** I am grateful to Maria Christakis, Sophia Drossopoulou, Peter Müller, and David Pichardie for comments on an earlier draft of this paper.

## References

1. Barnett, M., Chang, B.-Y.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A modular reusable verifier for object-oriented programs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) FMCO 2005. LNCS, vol. 4111, pp. 364–387. Springer, Heidelberg (2006)
2. Barnett, M., Fähndrich, M., Leino, K.R.M., Müller, P., Schulte, W., Venter, H.: Specification and verification: The Spec# experience. *Communications of the ACM* 54(6), 81–91 (2011)
3. Barrett, C., Tinelli, C.: CVC3. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 298–302. Springer, Heidelberg (2007)
4. Beckert, B., Hähnle, R., Schmitt, P.H. (eds.): *Verification of Object-Oriented Software*. LNCS (LNAI), vol. 4334. Springer, Heidelberg (2007)
5. Bertot, Y.: Filters on coinductive streams, an application to Eratosthenes’ sieve. In: Urzyczyn, P. (ed.) TLCA 2005. LNCS, vol. 3461, pp. 102–115. Springer, Heidelberg (2005)
6. Bertot, Y., Castéran, P.: *Interactive Theorem Proving and Program Development — Coq’Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. Springer (2004)
7. Bobot, F., Conchon, S., Contejean, E., Lescuyer, S.: Implementing polymorphism in SMT solvers. In: SMT 2008/BPR 2008: Proceedings of the Joint Workshops of the 6th International Workshop on Satisfiability Modulo Theories and 1st International Workshop on Bit-Precise Reasoning, pp. 1–5. ACM (July 2008)
8. Böhme, S., Nipkow, T.: Sledgehammer: Judgement day. In: Giesl, J., Hähnle, R. (eds.) IJCAR 2010. LNCS, vol. 6173, pp. 107–121. Springer, Heidelberg (2010)
9. Bove, A., Dybjer, P., Norell, U.: A brief overview of Agda — a functional language with dependent types. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLS 2009. LNCS, vol. 5674, pp. 73–78. Springer, Heidelberg (2009)
10. Bruttomesso, R., Pek, E., Sharygina, N., Tsitovich, A.: The OpenSMT solver. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 150–153. Springer, Heidelberg (2010)
11. 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: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLS 2009. LNCS, vol. 5674, pp. 23–42. Springer, Heidelberg (2009)
12. Cuoq, P., Kirchner, F., Kosmatov, N., Prevosto, V., Signoles, J., Yakobowski, B.: Frama-C — a software analysis perspective. In: Eleftherakis, G., Hinchey, M., Holcombe, M. (eds.) SEFM 2012. LNCS, vol. 7504, pp. 233–247. Springer, Heidelberg (2012)
13. de Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
14. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: A theorem prover for program checking. *Journal of the ACM* 52(3), 365–473 (2005)
15. Devillers, M., Griffioen, D., Müller, O.: Possibly infinite sequences in theorem provers: A comparative study. In: Gunter, E.L., Felty, A.P. (eds.) TPHOLS 1997. LNCS, vol. 1275, pp. 89–104. Springer, Heidelberg (1997)
16. Filliâtre, J.-C., Marché, C.: The Why/Krakatoa/Caduceus Platform for Deductive Program Verification. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 173–177. Springer, Heidelberg (2007)

17. Filliâtre, J.-C., Paskevich, A.: Why3 — where programs meet provers. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 125–128. Springer, Heidelberg (2013)
18. Guitton, J., Kanig, J., Moy, Y.: Why Hi-Lite Ada? In: Leino, K.R.M., Moskal, M. (eds.) BOOGIE 2011: First International Workshop on Intermediate Verification Languages (August 2011)
19. Harrison, J.: HOL Light: A tutorial introduction. In: Srivas, M., Camilleri, A. (eds.) FMCAD 1996. LNCS, vol. 1166, pp. 265–269. Springer, Heidelberg (1996)
20. Hoare, C.A.R.: An axiomatic basis for computer programming. *Communications of the ACM* 12(10), 576–580, 583 (1969)
21. Jacobs, B., Piessens, F.: The VeriFast program verifier. Technical Report CW-520, Department of Computer Science, Katholieke Universiteit Leuven (August 2008)
22. Jacobs, B., Smans, J., Piessens, F.: VeriFast: Imperative programs as proofs. In: VSTTE Workshop on Tools & Experiments (August 2010)
23. Kanig, J., Filliâtre, J.-C.: Who: A verifier for effectful higher-order programs. In: ACM SIGPLAN Workshop on ML. ACM (August 2009)
24. Kaufmann, M., Manolios, P., Moore, J.S.: *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers (2000)
25. Leclerc, F., Paulin-Mohring, C.: Programming with streams in Coq — A case study: The sieve of Eratosthenes. In: Barendregt, H., Nipkow, T. (eds.) TYPES 1993. LNCS, vol. 806, pp. 191–212. Springer, Heidelberg (1994)
26. Leino, K.R.M.: Dafny: An automatic program verifier for functional correctness. In: Clarke, E.M., Voronkov, A. (eds.) LPAR-16 2010. LNCS, vol. 6355, pp. 348–370. Springer, Heidelberg (2010)
27. Leino, K.R.M.: Automating induction with an SMT solver. In: Kuncak, V., Rybalchenko, A. (eds.) VMCAI 2012. LNCS, vol. 7148, pp. 315–331. Springer, Heidelberg (2012)
28. Leino, K.R.M., Moskal, M.: Usable auto-active verification. In: Ball, T., Zuck, L., Shankar, N. (eds.) UV 2010 (Usable Verification) Workshop (November 2010), <http://fm.csl.sri.com/UV10/>
29. Leino, K.R.M., Moskal, M.: Co-induction simply: Automatic co-inductive proofs in a program verifier. Technical Report MSR-TR-2013-49, Microsoft Research (May 2013)
30. Leino, K.R.M., Müller, P.: A basis for verifying multi-threaded programs. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 378–393. Springer, Heidelberg (2009)
31. Leino, K.R.M., Polikarpova, N.: Verified calculations. In: Fifth Working Conference on Verified Software: Theories, Tools, and Experiments, VSTTE 2013 (May 2013)
32. Milner, R.: *A Calculus of Communication Systems*. LNCS, vol. 92. Springer, Heidelberg (1980)
33. Nanevski, A., Morrisett, G., Shinnar, A., Govereau, P., Birkedal, L.: Ynot: Dependent types for imperative programs. In: Hook, J., Thiemann, P. (eds.) Proceeding of the 13th ACM SIGPLAN International Conference on Functional Programming, ICFP 2008, pp. 229–240. ACM (September 2008)
34. Nipkow, T., Paulson, L.C., Wenzel, M.T.: Isabelle/HOL. LNCS, vol. 2283. Springer, Heidelberg (2002)
35. Norell, U.: Towards a practical programming language based on dependent type theory. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology (September 2007)
36. Owre, S., Rushby, J.M., Shankar, N.: PVS: A prototype verification system. In: Kapur, D. (ed.) CADE 1992. LNCS (LNAI), vol. 607, pp. 748–752. Springer, Heidelberg (1992)
37. Pfenning, F., Schürmann, C.: System description: Twelf — A meta-logical framework for deductive systems. In: Ganzinger, H. (ed.) CADE 1999. LNCS (LNAI), vol. 1632, pp. 202–206. Springer, Heidelberg (1999)

38. Régis-Gianas, Y., Pottier, F.: A Hoare logic for call-by-value functional programs. In: Audebaud, P., Paulin-Mohring, C. (eds.) MPC 2008. LNCS, vol. 5133, pp. 305–335. Springer, Heidelberg (2008)
39. Reif, W.: The KIV system: Systematic construction of verified software. In: Kapur, D. (ed.) CADE 1992. LNCS, vol. 607, pp. 753–757. Springer, Heidelberg (1992)
40. Stump, A., Deters, M., Petcher, A., Schiller, T., Simpson, T.: Verified programming in Guru. In: Altenkirch, T., Millstein, T. (eds.) PLPV 009 — Proceedings of the 3rd Workshop on Programming Languages Meets Program Verification, pp. 49–58. ACM (January 2009)
41. Swamy, N., Chen, J., Fournet, C., Strub, P.-Y., Bhargavan, K., Yang, J.: Secure distributed programming with value-dependent types. In: Chakravarty, M.M.T., Hu, Z., Danvy, O. (eds.) Proceeding of the 16th ACM SIGPLAN International Conference on Functional Programming, ICFP 2011, pp. 266–278. ACM (September 2011)
42. Wenzel, M.: Isabelle/Isar—A versatile environment for human-readable formal proof documents. PhD thesis, Institut für Informatik, Technische Universität München (2002)