

Lecture Notes on Induction

Frank Pfenning

Carnegie Mellon University

Lecture 11

Thursday, March 11, 2021

1 Introduction

After analyzing first-order logic in the last lecture, we broach the subject of arithmetic and specifically induction in this lecture. In particular, we'll explore ways to carry out induction in Why3.

After a few introductory examples we explore the fundamentally inductive nature of α^* in dynamic logic—something we have avoided so far. This will allow us to prove the induction axiom for reasoning about programs, which we were unable to do before.

Learning goals. After this lecture, you should be able to:

- Use explicit induction in Why3;
- Reason about the semantic validity of axioms in dynamic logic using mathematical induction.

2 Peano's Axioms

The theory of arithmetic as axiomatized by Peano concerns the natural numbers $0, 1, 2, \dots$ while in Why3 we have been dealing with integers $\dots, -2, -1, 0, 1, 2, \dots$. This gap is not difficult to bridge, but in this lecture we will focus on natural numbers. You can read more about [Peano's Axioms](#) on Wikipedia.

He starts with axioms on equality, such as reflexivity, symmetry, and transitivity of equality. Furthermore, there is a constant 0 and successor function S such that $S(n) =$

$S(k)$ iff $n = k$. Furthermore, $0 \neq S(n)$ for all n . Finally, we have the induction axiom

$$P(0) \wedge (\forall n. P(n) \rightarrow P(S(n))) \rightarrow \forall n. P(n)$$

where $P(n)$ is an arbitrary predicate on natural numbers. In the context of this course, $P(n)$ is a property of n definable by a formula.

We can also write the axiom as an inference rule in the sequent calculus.

$$\frac{\Gamma \vdash P(0), \Delta \quad \Gamma, P(n) \vdash P(S(n)), \Delta}{\Gamma \vdash P(e), \Delta} \text{ ind}$$

where n does not already occur in Γ , $P(e)$, or Δ . This rule is formulated in this particular manner because we like to avoid mentioning other connectives or quantifiers and keep the rules “pure”.

3 Induction in Why3

Among the Why3 standard libraries we find `int.SimpleInduction` for mathematical induction and also `int.Induction` for complete induction.

```

1 module SimpleInduction
2
3   use Int
4
5   predicate p int
6
7   axiom base: p 0
8
9   axiom induction_step: forall n:int. 0 <= n -> p n -> p (n+1)
10
11  lemma SimpleInduction : forall n:int. 0 <= n -> p n
12
13 end

```

This module keeps the predicate p as well as the base and induction_step axioms *abstract* and provides the lemma `SimpleInduction`. The idea is to *clone* this module by providing a concrete predicate for p . Secondly, we want to turn base and induction_step into *lemmas*, which creates corresponding proof obligations. If they can be proved, we obtain those two lemmas, plus `SimpleInduction` as a consequence.

Here is an example. We would like to prove

$$\sum_{i=0}^{i=n} i = \frac{n(n+1)}{2}$$

We start by defining a sum function with a lower bound of a and upper bound of b , inclusive.

```

1 let rec function sum (a : int) (b : int) : int =
2   variant { b - a }
3   if a > b then 0 else sum a (b-1) + b

```

Then we define the predicate $P(n)$ intended for the induction, here called `sum_square`.

```
1 predicate sum_square (n : int) =
2   n >= 0 -> sum 0 n = div (n*(n+1)) 2
```

We can now clone the standard library as sketched above.

```
1 clone int.SimpleInduction
2   with predicate p = sum_square, lemma base, lemma induction_step
```

The summary of the whole theory is below. Just to be sure, we also restate the desired property as a *goal*. This generates a proof obligation just like a lemma, but does not assume the proven formula. This is helpful if we'd like to avoid polluting the search space.

```
1 theory SumSquare1
2
3   use int.Int
4   use int.EuclideanDivision
5
6   let rec function sum (a : int) (b : int) : int =
7     variant { b - a }
8     if a > b then 0 else sum a (b-1) + b
9
10  predicate sum_square (n : int) =
11    n >= 0 -> sum 0 n = div (n*(n+1)) 2
12
13  clone int.SimpleInduction
14    with predicate p = sum_square, lemma base, lemma induction_step
15
16  goal G : forall n:int. n >= 0 -> sum_square n
17
18 end
```

4 Induction via Recursion

Besides the technique above we can also often provide an explicit witness as a function. When we can prove the termination of such a witness function, and the contract expresses the desired property, then the theorem we would like to state becomes an immediate consequence.

Here is a first way of achieving this¹.

```
1 theory SumSquare3
2
3   use int.Int
4   use int.EuclideanDivision
5
6   let rec function sum (a : int) (b : int) : int =
7     variant { b - a }
8     if a > b then 0 else sum a (b-1) + b
9
```

¹which we did not do in lecture

```

10 let rec function sum_rec (n : int) =
11   requires { n >= 0 }
12   variant { n }
13   ensures { result = div (n * (n+1)) 2 }
14   ensures { result = sum 0 n }
15   if n = 0 then 0
16   else sum_rec (n-1) + n
17
18   goal G : forall n:int. n >= 0 -> sum 0 n = div (n * (n+1)) 2
19
20 end

```

The witness function here is `sum_rec`, with an explicit property statement following. Essentially, `sum_rec` is a *circular proof* of the desired property, acceptable because the variant shows that it is terminating.

An even slicker formulation of this technique² is for the witness function to return a Boolean, in effect making it a lemma. The key here is that the postcondition must guarantee the property we ultimately want, so it might use the same predicate we used with the explicit induction axiom.

```

1 theory SumSquare2
2
3   use int.Int
4   use int.EuclideanDivision
5
6   let rec function sum (a : int) (b : int) : int =
7     variant { b - a }
8     if a > b then 0 else sum a (b-1) + b
9
10  predicate sum_square (n : int) =
11    n >= 0 -> sum 0 n = div (n * (n+1)) 2
12
13  let rec lemma sq (n : int) =
14    requires { n >= 0 }
15    variant { n }
16    ensures { result /\ sum_square n }
17    if n = 0 then true
18    else sq (n-1)
19
20  goal G : forall n:int. n >= 0 -> sum_square n
21
22 end

```

The verification condition in the $n = 0$ branch is exactly the base of the induction, and the condition in the $n > 0$ branch is the induction step. In the tail call `sq (n - 1)` we have to satisfy the precondition ($n \geq 0$) which is true, and then get to assume the postcondition on the actual parameter, that is, `sum_square (n - 1)`. We have to prove `sum_square n`, precisely the implication that represents the induction step.

²which we *did* briefly show in lecture

5 Revisiting Dynamic Logic in Why3

When we gave the mathematical definition of repetition in Dynamic logic we wrote

$$\omega_0 \llbracket \alpha^* \rrbracket \omega_n \quad \text{iff there exist } \omega_1, \dots, \omega_{n-1} \text{ such that } \omega_i \llbracket \alpha \rrbracket \omega_{i+1} \text{ for all } 0 \leq i < n.$$

In our formalization we instead used the axiom

```
1  axiom run_star : forall omega alpha nu.
2    run omega (Star alpha) nu
3    <-> omega = nu \/ exists mu. run omega alpha mu /\ run mu (Star
      alpha) nu
```

This formalization is considerably weaker since it does not allow us to prove properties of α^* by induction. In particular, we are unable to prove the dynamic logic axiom of induction:

```
1  axiom induction : forall omega alpha q.
2    models omega (Box (Star alpha) q)
3    <-> models omega (And q (Box (Star alpha)
4      (Implies q (Box alpha q))))
```

Our mathematical argument for this property relied on induction on n as used in the mathematical definition of $\omega_0 \llbracket \alpha^* \rrbracket \omega_n$.

Now that we have discovered how to carry out inductive arguments in Why3 we can revise our definition to match the mathematical definition more closely and then use that to actually prove the induction axiom. The complete code can be found in [ndl.mlw](#).

As a first step, we need two predicates: `run` as before and `run_bdd` that a relation running a program α for n times in succession, written as $\omega \llbracket \alpha \rrbracket^n \nu$. Mathematically we define

$$\begin{aligned} \omega \llbracket \alpha^* \rrbracket \nu & \quad \text{iff there exists an } n \text{ such that } \omega \llbracket \alpha \rrbracket^n \nu \\ \omega \llbracket \alpha \rrbracket^0 \nu & \quad \text{iff } \omega = \nu \\ \omega \llbracket \alpha \rrbracket^{n+1} \nu & \quad \text{iff there exists a } \mu \text{ such that } \omega \llbracket \alpha \rrbracket \mu \text{ and } \mu \llbracket \alpha \rrbracket^n \nu \end{aligned}$$

Our formalization in Why3 models the mathematical definition very closely.

```
1  predicate run (omega : state) (alpha : prog) (nu : state)
2  predicate run_bdd (omega : state) (alpha : prog) (n : int) (nu :
      state)
3
4  axiom run_star : forall omega alpha nu.
5    run omega (Star alpha) nu
6    <-> exists n. n >= 0 /\ run_bdd omega alpha n nu
7  axiom run_bdd : forall omega alpha nu n.
8    run_bdd omega alpha n nu
9    <-> (n = 0 /\ omega = nu)
10     \/ (n > 0 /\ exists mu. run omega alpha mu
11       /\ run_bdd mu alpha (n-1) nu)
```

We elide the remaining axioms for assignment, sequential composition, nondeterministic choice, and guards because they have not changed from [Lecture 8](#).

Why3 can now prove the original axiom, which corresponds to unfolding the repetition (implication from left to right) and folding the repetition (from right to left). Now induction is required here.

```

1 lemma run_star_fold : forall omega alpha nu.
2   (omega = nu \/ exists mu. run omega alpha mu
3     /\ run mu (Star alpha) nu)
4   -> run omega (Star alpha) nu
5
6 lemma run_star_unfold : forall omega alpha nu.
7   run omega (Star alpha) nu
8   -> (omega = nu \/ exists mu. run omega alpha mu
9     /\ run mu (Star alpha) nu)

```

Similarly, the loop unfolding axiom of dynamic logic, namely

$$[\alpha^*]Q \leftrightarrow Q \wedge [\alpha][\alpha^*]Q$$

is easy to prove in both directions.

```

1 lemma models_box_star : forall omega alpha q.
2   models omega (Box (Star alpha) q)
3   <-> models omega (And q (Box alpha (Box (Star alpha) q)))

```

As before, we have mapped the bi-implication from dynamic logic directly to bi-implication in Why3.

What remains now is to prove the representation of the induction axiom in Why3, which was impossible before.

6 Proving Induction for Dynamic Logic in Why3

We would like to prove the validity of computational induction in dynamic logic:

$$[\alpha^*]Q \leftrightarrow Q \wedge [\alpha^*](Q \rightarrow [\alpha]Q)$$

In our representation both directions of this bi-implication require induction.

We start by proving the left-to-right direction using several lemmas. We first state them in ordinary mathematical language

*For all $n \geq 0$, ω , and ν ,
 if $\omega \llbracket \alpha \rrbracket^n \mu$ and $\mu \llbracket \alpha \rrbracket \nu$ for some μ
 then $\omega \llbracket \alpha \rrbracket \mu'$ and $\mu' \llbracket \alpha \rrbracket^n \nu$ for some μ'*

We prove this by induction on n . In Why3:

```

1 predicate run_bdd_left (n:int) =
2   forall omega alpha nu.
3     (exists mu. run_bdd omega alpha n mu /\ run mu alpha nu)
4     -> (exists mu'. run omega alpha mu' /\ run_bdd mu' alpha n nu)
5
6 clone int.SimpleInduction as Bdd1
7 with predicate p = run_bdd_left, lemma base, lemma induction_step

```

Note that we clone `int.SimpleInduction` with an explicit name `Bdd1` so that we can clone the same module later without creating a name conflict. We can similarly prove the opposite direction of this implication.

```

1  predicate run_bdd_right (n:int) =
2    forall omega alpha nu.
3      (exists mu'. run omega alpha mu' /\ run_bdd mu' alpha n nu)
4      -> (exists mu. run_bdd omega alpha n mu /\ run mu alpha nu)
5
6  clone int.SimpleInduction as Bdd2
7    with predicate p = run_bdd_right, lemma base, lemma
      induction_step

```

With these, we can prove an alternative characterization of n -fold repetition, namely a “right-associative” composition instead of the original “left-associative” one.

$$\begin{aligned} \omega[[\alpha]]^0\nu & \text{ iff } \omega = \nu \\ \omega[[\alpha]]^{n+1}\nu & \text{ iff there exists a } \mu \text{ such that } \omega[[\alpha]]^n\mu \text{ and } \mu[[\alpha]]\nu \end{aligned}$$

This can be proved directly from the two preceding lemmas.

```

1  lemma run_bdd_alt : forall omega alpha nu n.
2    run_bdd omega alpha n nu
3    <-> (n = 0 /\ omega = nu)
4        \/ (n > 0 /\ exists mu. run_bdd omega alpha (n-1) mu /\ run mu
5          alpha nu)
6
7  lemma box_star_0 : forall omega alpha q.
8    models omega (Box (Star alpha) q) -> models omega q
9
10 lemma box_star_left : forall omega alpha q.
11   models omega (Box (Star alpha) q)
12   -> models omega (Box (Star alpha) (Box alpha q))
13
14 lemma induction_unwind : forall omega alpha q.
15   models omega (Box (Star alpha) q)
16   -> models omega (And q (Box (Star alpha)
17     (Implies q (Box alpha q))))

```

The other direction of the induction axiom is intuitively more complicated, but follows here just by one lemma, proved with a simple induction.

```

1  predicate run_bdd_inv (n : int) =
2    forall omega alpha q.
3      models omega q /\ models omega (Box (Star alpha)
4        (Implies q (Box alpha q)))
5      -> forall nu. run_bdd omega alpha n nu -> models nu q
6
7  clone int.SimpleInduction as Bdd3
8    with predicate p = run_bdd_inv, lemma base, lemma induction_step
9
10 lemma induction_wind : forall omega alpha q.
11   models omega (And q (Box (Star alpha) (Implies q (Box alpha q))))
12   -> models omega (Box (Star alpha) q)

```

7 Induction with Invariants

Induction with invariants just requires a few lemmas regarding validity (as expressed with $\Box P$) and the properties we have already proved. We just summarize them here, since they introduce no new and interesting ideas.

```

1  lemma valid_box : forall omega alpha p.
2    models omega (Valid p) -> models omega (Box alpha p)
3
4  lemma dist_box_implies : forall omega alpha p q.
5    models omega (Box alpha (Implies p q))
6    -> models omega (Box alpha p)
7    -> models omega (Box alpha q)
8
9  lemma dist_valid_box : forall omega alpha p q.
10   models omega (Valid (Implies p q))
11   -> models omega (Box alpha p)
12   -> models omega (Box alpha q)
13
14  lemma induction_valid : forall omega q alpha.
15   models omega (And q (Valid (Implies q (Box alpha q))))
16   -> models omega (Box (Star alpha) q)
17
18  lemma induction_inv : forall omega j alpha q.
19   models omega (And j
20                 (And (Valid (Implies j (Box alpha j)))
21                     (Valid (Implies j q))))
22   -> models omega (Box (Star alpha) q)

```

The final theorem here is what we were aiming at, namely

$$J \wedge \Box(J \rightarrow [\alpha]J) \wedge \Box(J \rightarrow Q) \rightarrow [\alpha^*]Q$$

What we did not accomplish and still remains to be done is the formalization of the axiom of *convergence* which concerns $\langle \alpha^* \rangle Q$.