

Lecture Notes on Induction

Ruben Martins*

Carnegie Mellon University

Lecture 9

Tuesday, February 13, 2024

1 Introduction

In [Lecture 7](#) we analyzed reasoning about while loops. To do so, we introduced the more abstract nondeterministic dynamic logic (NDL) with nondeterministic choice and iteration. Eventually, we arrived at two axioms for iteration in NDL: unrolling and induction:

$$\begin{aligned} [\alpha^*]Q &\leftrightarrow Q \wedge [\alpha][\alpha^*]Q && \text{unrolling} \\ [\alpha^*]Q &\leftrightarrow Q \wedge [\alpha^*](Q \rightarrow [\alpha]Q) && \text{induction} \end{aligned}$$

From these and the other axioms, we were able to derive a *practical* version of the axiom for reasoning about while loops with loop invariants J . This is no longer an “if and only if” but nevertheless allows us to prove properties of while loops.

$$[\text{while } P \ \alpha]Q \leftarrow J \wedge \Box(J \wedge P \rightarrow [\alpha]J) \wedge \Box(J \wedge \neg P \rightarrow Q)$$

Recall that $\Box P$ is true if P is valid, which means true in every possible state.

This final axiom is precisely the form that Why3 (and other systems) use to generate verification conditions for loops when provided with a loop invariant J . We still reason about *partial correctness*, that is, this axiom only ensures that *if the program terminates* then the postcondition Q will hold in the final state. In the next lecture we will address *total correctness*, which also requires termination.

Our formalization of NDL in Why3 in [Lecture 8](#) proved correctness for various axioms allowing decomposition of α in $[\alpha]Q$, but omitted NDL axioms for assignment

*Closely adapted from notes written by Frank Pfenning in Spring 2022

and iteration. The first one is possible, but tedious due to the side condition on occurrences of variables; the second one is the subject of this lecture. So our goal is to prove the correctness of the induction axiom in NDL.

Perhaps not surprisingly, this seems to require proof by *induction*. In general proofs by induction are difficult to automate, so they are not natively available in Why3 or the back-end provers such as alt-ergo, Z3, or CVC4. Instead we have to “trick” Why3 into carrying out such proofs for us. We show how to do this, and then apply the learned techniques to the induction axiom in NDL.

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.

3 Induction Proofs in NDL

So far, we have been using NDL and its axioms to turn reasoning about programs into problems of logical reasoning. Now we turn the tables: we want to see how to model logical reasoning by program reasoning! The goal, in short, is to write a program α so that $[\alpha]P(i)$ is true if there is an induction proof $\forall i. P(i)$. Actually, we want something slightly stricter, namely to encoding particular induction schemas. In the first case, it will simply Peano’s axiom.

So consider the following proposition in NDL:

$$[i \leftarrow 0 ; (i \leftarrow i + 1)^*]P(i)$$

Because the loop may be executed any number of times $0, 1, 2, \dots$, this proposition is valid in NDL if $P(i)$ holds for all $i \geq 0$. Let's reason through a verification effort using our axioms for NDL:

$$\begin{aligned}
 & \models [i \leftarrow 0 ; (i \leftarrow i + 1)^*]P(i) \\
 \text{iff } & \models [i \leftarrow 0][i \leftarrow i + 1]^*P(i) \\
 \text{iff } & \models \forall i. i = 0 \rightarrow [(i \leftarrow i + 1)^*]P(i) \\
 \text{if } & \models \forall i. i = 0 \rightarrow P(i) \wedge \Box(P(i) \rightarrow [i \leftarrow i + 1]P(i))
 \end{aligned}$$

Here, the last step used the form of the induction axiom where we have replaced $[\alpha^*]Q$ by $\Box Q$. We can take some further steps to eliminate the \Box modality.

$$\begin{aligned}
 & \models \forall i. i = 0 \rightarrow P(i) \wedge \Box(P(i) \rightarrow [i \leftarrow i + 1]P(i)) \\
 \text{iff } & \models \forall i. i = 0 \rightarrow P(i) \wedge \Box(P(i) \rightarrow \forall i'. i' = i + 1 \rightarrow P(i')) \\
 \text{iff } & \models (\forall i. i = 0 \rightarrow P(i)) \wedge \forall i. i = 0 \rightarrow \Box(P(i) \rightarrow \forall i'. i' = i + 1 \rightarrow P(i')) \\
 \text{iff } & \models P(0) \wedge \forall i. i = 0 \rightarrow \Box(P(i) \rightarrow P(i + 1))
 \end{aligned}$$

In the second conjunct, we lose the assumption about $i = 0$ underneath the \Box modality because we have to prove that $P(i) \rightarrow P(i + 1)$ regardless of the state. We can formulate the upshot of this line of reasoning as

$$\models [i \leftarrow 0 ; (i \leftarrow i + 1)^*]P(i) \text{ if } \models P(0) \wedge \forall i. P(i) \rightarrow P(i + 1)$$

In other words, using the induction axiom in NDL, we can reduce the proof of proving the validity of $[i \leftarrow 0 ; (i \leftarrow i + 1)^*]P(i)$ to proving the validity of $P(0) \wedge \forall i. P(i) \rightarrow P(i + 1)$, which should have been intuitively clear when we started.

4 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
```

5 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 this⁴.

```

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
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 technique⁵ 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

```

⁴which we did not do in lecture

⁵which we *did* briefly show in lecture

```

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 $\text{sq}(n-1)$ 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, $\text{sum_square}(n-1)$. We have to prove $\text{sum_square } n$, precisely the implication that represents the induction step.

6 Revisiting Dynamic Logic in Why3

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

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

We now continue our formalization of NDL in Why3. You can find the code in the file [ndl-v2.mlw](#).

In order to represent the above definition, we define a new auxiliary predicate `run_bdd` and give the definition of $\llbracket \alpha^* \rrbracket$ in terms of the new predicate.

```

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

```

The formalization of `run_bdd` is via two simple axioms.

```

1  axiom run_bdd_0 : forall omega alpha nu.
2    run_bdd omega alpha 0 nu <-> omega = nu
3  axiom run_bdd_1 : forall omega alpha n nu. n > 0 ->
4    run_bdd omega alpha n nu
5    <-> exists mu. run omega alpha mu /\ run_bdd mu alpha (n-1) nu

```

We'd now like to show that this implies the recursive specification, namely that

$$\omega \llbracket \alpha^* \rrbracket \nu \text{ iff } \omega = \nu \vee \exists \mu. \omega \llbracket \alpha \rrbracket \mu \wedge \mu \llbracket \alpha^* \rrbracket \nu$$

Unfortunately, Why3 could not establish this directly, but we needed two simple lemmas. We conjecture that the existential quantification over $n \geq 0$ in the definition of $\llbracket \alpha^* \rrbracket$ is to blame. In any case, with those two lemmas we can prove the above characterization.

```

1 lemma run_bdd_left0 : forall omega alpha.
2   run omega (Star alpha) omega
3
4 lemma run_bdd_left1 : forall omega alpha mu n nu. n >= 0 ->
5   run omega alpha mu /\ run_bdd mu alpha n nu
6   -> run_bdd omega alpha (n+1) nu
7
8 lemma run_star_rec : forall omega alpha nu.
9   (omega = nu
10    /\ exists mu. run omega alpha mu /\ run mu (Star alpha) nu)
11   <-> run omega (Star alpha) nu

```

Similarly, the loop unrolling 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 q /\ models omega (Box alpha (Box (Star alpha) q))

```

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

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

7 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 $\omega \Vdash \alpha^n \mu$ and $\mu \Vdash \alpha \nu$ for some μ
 then $\omega \Vdash \alpha \mu'$ and $\mu' \Vdash \alpha^n \nu$ for some μ'*

We prove this by induction on n . In Why3:

```

1 predicate run_bdd_left (n:int) = forall omega alpha nu.
2   (exists mu. run_bdd omega alpha n mu /\ run mu alpha nu)
3   -> (exists mu'. run omega alpha mu' /\ run_bdd mu' alpha n nu)
4
5 let rec lemma ind_run_bdd_left (n:int) : unit =
6   requires { n >= 0 }
7   variant { n }
8   ensures { run_bdd_left n }
9   if n = 0 then () else ind_run_bdd_left (n-1)

```

We can similarly prove the opposite direction of this implication, but it does not appear to be necessary: We can already prove one direction of the induction axiom without any further induction.

```
1 lemma induction_unwind : forall omega alpha q.
2   models omega (Box (Star alpha) q)
3   -> models omega q
4   /\ models omega (Box (Star alpha) (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) = forall omega alpha q.
2   models omega q
3   /\ models omega (Box (Star alpha) (Implies q (Box alpha q)))
4   -> forall nu. run_bdd omega alpha n nu -> models nu q
5
6 let rec lemma ind_run_bdd_inv (n:int) : unit =
7   requires { n >= 0 }
8   variant { n }
9   ensures { run_bdd_inv n }
10  if n = 0 then () else ind_run_bdd_inv (n-1)
11
12 lemma induction_wind : forall omega alpha q.
13   models omega q
14   /\ models omega (Box (Star alpha) (Implies q (Box alpha q)))
15   -> models omega (Box (Star alpha) q)
```

8 Induction with Invariants

Induction with invariants just requires one lemma regarding validity (as expressed with $\Box P$) and the properties we have already proved. We just summarize the lemma and two theorems here. These can be found at the end of the file [ndl-v2.mlw](#).

```
1 (* lemma: Q /\ [] (Q -> [alpha]Q) -> [alpha]*Q *)
2 lemma induction_valid : forall omega q alpha.
3   models omega (And q (Valid (Implies q (Box alpha q))))
4   -> models omega (Box (Star alpha) q)
5
6 (* lemma: [] (P -> Q) -> [alpha]P -> [alpha]Q *)
7 lemma dist_valid_box : forall omega alpha p q.
8   models omega (Valid (Implies p q))
9   -> models omega (Box alpha p)
10  -> models omega (Box alpha q)
11
12 (* induction with invariant J *)
13 lemma induction_inv : forall omega j alpha q.
14   models omega (And j
15     (And (Valid (Implies j (Box alpha j)))
16         (Valid (Implies j q))))
17   -> 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 $\langle \alpha^* \rangle Q$.