

Coq projects for dIFP 2020

Robbert Krebbers

October 29, 2020

This file gives a list of suggestions for the Coq project of dIFP. You may also invent a project yourself or do a variation of one of the projects. If you invent a project yourself or do a variation of one of the suggested projects, you have to discuss it with the teacher.

Library file. As part of your Coq project you should make use of the following library file:

`http://www.cs.au.dk/~spitters/project_lib.v`

This file contains the definition of the `Case` tactics, as well as some common Coq definitions, such as identifiers.

Extensions. Each project ends with an exercise that lists possible extensions, of which you have to choose at least one. You are advised to first complete the whole exercise before starting to work on the extension. Obviously, completing more difficult extensions is necessary for a higher grade. For a top grade, you are expected to complete a very difficult extension.

In the report, please follow the structure of the exercises, so that it is clear which extensions you have completed.

1 The simply typed λ -calculus

The simply typed λ -calculus is a subsystem of Coq that only has function types. The goal of this project is to formalize the simply typed λ -calculus in Coq and to implement a type checker for it.

The types and terms are given by the following grammar:

$$\begin{aligned} A, B &::= X \mid A \rightarrow B \\ t, r &::= x \mid tr \mid \lambda x : A . t \end{aligned}$$

Here, X ranges over a set of type variables and x over a set of term variables. Examples of types are $X \rightarrow Y \rightarrow X$ and $(X \rightarrow X) \rightarrow (X \rightarrow X)$, and examples of terms are $\lambda x : X . \lambda y : Y . x$ and $\lambda f : X \rightarrow X . \lambda x : X . f (f x)$.

In order to define the set of well-typed terms, we define a typing judgment $\Gamma \vdash t : A$ that states that *a term t has type A in environment Γ* . The environment Γ associates types to term variables. The typing judgment $\Gamma \vdash t : A$ is defined by the following inference rules:

$$\frac{x : A \in \Gamma}{\Gamma \vdash x : A} \quad \frac{\Gamma \vdash t : A \rightarrow B \quad \Gamma \vdash r : A}{\Gamma \vdash tr : B} \quad \frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x : A . t : A \rightarrow B}$$

For example, the term $\lambda x : X . \lambda y : Y . x$ has type $X \rightarrow Y \rightarrow X$, which is shown by the following derivation:

$$\frac{\frac{x : X, y : Y \vdash x : X}{x : X \vdash \lambda y : Y . x : Y \rightarrow X}}{\vdash \lambda x : X . \lambda y : Y . x : X \rightarrow Y \rightarrow X}$$

We formalize the types and terms of the simply typed λ -calculus using the following inductive definitions in Coq:

```

Inductive type :=
  | tvar : id -> type
  | tarr : type -> type -> type.

Inductive term :=
  | var : id -> term
  | app : term -> term -> term
  | lam : id -> type -> term -> term.

```

The type `id` of identifiers has been introduced in the chapter “Imp” of Software Foundations and is defined in the supplemented library file.

Exercise 1.1 Give Coq definitions of type `type` corresponding to:

$$X \rightarrow Y \rightarrow X \quad \text{and} \quad (X \rightarrow X) \rightarrow (X \rightarrow X).$$

Exercise 1.2 Give Coq definitions of type `term` corresponding to:

$$\lambda x : X . \lambda y : Y . x \quad \text{and} \quad \lambda f : X \rightarrow X . \lambda x : X . f (f x).$$

We will represent environments Γ as partial functions:

```

Definition env := id -> option type.
Definition empty_env : env := fun x => None.
Definition override (E : env) (x : id) (A : type) : env :=
  fun y => if beq_id x y then Some A else E y.

```

Exercise 1.3 Represent the typing judgment $\Gamma \vdash t : A$ as an inductively defined proposition. Do so by completing the following definition:

```

Inductive typed : env -> term -> type -> Prop :=
  | var_typed : forall E x A,
    E x = Some A ->
    typed E (var x) A
  (* fill out *).

```

Exercise 1.4 Prove $\vdash \lambda x : X . \lambda y : Y . x : X \rightarrow Y \rightarrow X$. You have to create a lemma of the following shape:

```

Lemma test : typed (*your environment*) (*term*) (*type*).

```

Exercise 1.5 Implement a type checker by filling out the following definition:

```
Fixpoint typecheck (E : env) (t : term) : option type :=
  (* fill out *).
```

Exercise 1.6 Explain the difference between the typing judgment `typed` and the type checker `typecheck`.

Exercise 1.7 Write 2 positive tests and 2 negative tests of the type checker and prove these using the `reflexivity` tactic.

Exercise 1.8 Prove completeness and soundness of your type checker with respect to the typing rules:

```
Lemma typecheck_complete : forall E t A,
  typed E t A ->
  typecheck E t = Some A.
Lemma typecheck_sound : forall E t A,
  typecheck E t = Some A ->
  typed E t A.
```

To prove completeness, you may want to use `induction` on the hypothesis `typed E t A`. To prove soundness, you may want to vary the induction hypothesis.

Exercise 1.9 Extend the language with a feature of your choice. For example:

- Products/conjunctions and sums/disjunctions (*easy*).
- Natural numbers and recursion (*moderate*).
- Universal/existential quantification (*difficult*).
- A relation for β -reduction and a proof of type preservation (*very difficult*). See https://en.wikipedia.org/wiki/Subject_reduction.

You have to extend the types, terms, typing rules, type checker and proofs.

2 SAT solver

The goal of this project is to implement a Boolean Satisfiability (SAT) solver in Coq and to prove soundness of your implementation.

Boolean formulas are given by the following grammar:

$$p, q ::= x \mid \text{true} \mid \text{false} \mid p \wedge q \mid p \vee q \mid p \rightarrow q \mid \neg p.$$

Here, x ranges over an infinite set of propositional variables. Examples of formulas are $(x \vee \neg y) \wedge (\neg x \vee y)$, $y \rightarrow (x \vee y)$ and $x \wedge \neg x \wedge \text{true}$.

A Boolean formula p is said to be *satisfiable* in case its truth table contains at least one row whose outcome is 1. For example:

x	y	$(x \vee \neg y) \wedge (\neg x \vee y)$	x	y	$\neg y \rightarrow (x \vee y)$	x	$x \wedge \neg x \wedge \text{true}$
0	0	1	0	0	0	0	0
0	1	0	0	1	1	1	0
1	0	0	1	0	1		<i>unsatisfiable</i>
1	1	1	1	1	1		
<i>satisfiable</i>			<i>satisfiable</i>				

The first part of this project is to represent Boolean formulas in Coq and to formally define the notion of satisfiability.

Exercise 2.1 Define an inductive type form that represents Boolean formulas. Start by completing the following definition:

```

Inductive form :=
  | var : id -> form
  (* fill out *).
```

The type `id` of identifiers has been introduced in the chapter “Imp” of Software Foundations and is defined in the supplemented library file.

Exercise 2.2 Give Coq definitions of type `form` corresponding to:

$$(x \vee \neg y) \wedge (\neg x \vee y) \quad \neg y \rightarrow (x \vee y) \quad \text{and} \quad x \wedge \neg x \wedge \text{true}.$$

In order to define satisfiability we introduce the notion of a *valuation*, which is a function that assigns true or false to each propositional variable.

```

Definition valuation := id -> bool.
Definition empty_valuation : valuation := fun x => false.
Definition override (V : valuation)
  (x : id) (b : bool) : valuation :=
  fun y => if beq_id x y then b else V y.
```

Exercise 2.3 Define the interpretation of a formula (using the ‘truth table semantics’) by filling out the following definition:

```

Fixpoint interp (V : valuation) (p : form) : bool :=
  (* fill out *).
```

A formula is said to be *satisfiable* if there exists a valuation that makes the formula true. This corresponds to the following definition in Coq:

```

Definition satisfiable (p : form) : Prop :=
  exists V : valuation, interp V p = true.
```

You may want to read the section “Existential Quantification” of the chapter “MoreLogic” of Software Foundations.

Exercise 2.4 Prove in Coq that $(x \vee \neg y) \wedge (\neg x \vee y)$ and $\neg y \rightarrow (x \vee y)$ are satisfiable. You should create two lemmas of the shape below:

```
Lemma test1 : satisfiable (*formula 1*).
Lemma test2 : satisfiable (*formula 2*).
```

Exercise 2.5 Define a function that given a formula computes a valuation in which the formula is true. You should implement this function by enumerating all possible valuations (that is generating its truth table).

```
Definition find_valuation (p : form) : option valuation :=
  (* fill out *).
```

The function `find_valuation` should yield `None` in case no such valuation exists (i.e. the formula is unsatisfiable).

We can now define our SAT solver as follows:

```
Definition solver (p : form) : bool :=
  match find_valuation p with
  | Some _ => true
  | None => false
  end.
```

Exercise 2.6 Explain the difference between `satisfiable` and `solver`.

Exercise 2.7 Write 2 positive and 2 negative tests of the solver and prove these tests using the `reflexivity` tactic.

Exercise 2.8 Prove that `solver` is sound. That means:

```
Lemma solver_sound : forall p,
  solver p = true -> satisfiable p.
```

Exercise 2.9 Extend your SAT solver in one of the following ways:

- Extend your development to three-valued logic (**moderate**). Three-valued logic has three truth values: `true`, `false` and some indeterminate third value. See https://en.wikipedia.org/wiki/Three-valued_logic.
- Write an optimizer that simplifies a Boolean formula using the laws below, and prove correctness of your optimizer (**moderate**).

$x \wedge \text{true} = x$	$\text{true} \wedge x = x$	$x \wedge \text{false} = \text{false}$	$\text{false} \wedge x = \text{false}$
$x \vee \text{true} = \text{true}$	$\text{true} \vee x = \text{true}$	$x \vee \text{false} = x$	$\text{false} \vee x = x$

You should incorporate this optimizer in `solver` and its correctness proof.

- Write a converter that transforms Boolean formulas into negation normal form and prove correctness of your converter (**difficult**). A Boolean formula is in negation normal form if the negation operator \neg is only applied to variables. You should incorporate this optimizer in `solver` and its correctness proof.
- Write a converter that transforms Boolean formulas into conjunctive normal form and prove correctness of your converter (**very difficult**). See also https://en.wikipedia.org/wiki/Conjunctive_normal_form. You should incorporate this optimizer in `solver` and its correctness proof.
- Prove completeness of your solver (**very difficult**). That means:

```
Lemma solver_complete : forall p,
  satisfiable p -> solver p = true.
```


3 Binary search trees

The goal of this project is to implement some operations on binary search trees in Coq and to prove essential properties of the implementation.

A binary search tree is a data structure to efficiently store a finite set of natural numbers. Binary trees are inductively defined in Coq as follows:

```
Inductive tree :=
  | leaf : tree
  | node : tree -> nat -> tree -> tree.
```

A binary tree is said to be a binary search tree if it satisfies the *binary search tree property*, which states that the value in each node is greater (but not equal to) than all values in the left sub-tree, and smaller (but not equal to) than all values in the right sub-tree. For example, the tree on the left satisfies the property, but the one on the right does not (leaves are omitted):



The tree on the right does not satisfy the binary search tree property because there is a node numbered 2 as a left sub-tree of a node numbered 1.

Exercise 3.1 Give Coq definitions of type `tree` corresponding to the above binary trees.

Exercise 3.2 Define an inductively defined proposition that describes whether a binary tree satisfies the binary search tree property. Do so by completing the following definition:

```
Inductive sorted : tree -> Prop :=
```

```
| leaf_sorted : sorted leaf
(* fill out *).
```

In order to define the above function you may want to define helpers `greater : nat -> tree -> Prop` and `smaller : nat -> tree -> Prop` or `all : (nat -> Prop) -> tree -> Prop` as inductively defined propositions.

Exercise 3.3 Prove that the first tree you have defined in Exercise 3.1 satisfies the binary search tree property, and that the second one does not. You may want to make use of composition ; of tactics and the `repeat` tactical.

Exercise 3.4 Define a function that describes whether an element is contained in a binary search tree or not:

```
Fixpoint elem_of (x : nat) (t : tree) : bool :=
(* fill out *).
```

To define an efficient implementation of the `elem_of` function, you should not traverse the whole tree, but make use of the binary search tree property.

Exercise 3.5 Write 2 positive tests and 2 negative tests of the `elem_of` function and prove these using the `reflexivity` tactic.

The last part of this project is to define a function that inserts a natural number into a binary search tree. You have to prove that the resulting tree is indeed a binary search tree, and you have to prove that it is correct.

Exercise 3.6 Define a function that inserts a natural number into a binary search tree:

```
Fixpoint insert (x : nat) (t : tree) : tree :=
(* fill out *).
```

If the number `x` is already in the tree, it should return the original tree.

Exercise 3.7 Write 2 tests of the `insert` function and prove these using the `reflexivity` tactic.

Exercise 3.8 Prove that the `insert` function preserves the binary search tree property. That means:

```
Lemma insert_sorted : forall t x,
sorted t -> sorted (insert x t).
```

Exercise 3.9 Prove that the `insert` function is correct. That means:

```
Lemma insert_correct : forall t x y,
  sorted t ->
  elem_of y (insert x t) = orb (elem_of y t) (beq_nat x y).
```

Exercise 3.10 Extend your development on binary search trees with a feature of your choice. For example:

- A converter from lists to binary search trees and vice versa. Prove a useful property about the interaction between these functions (*easy*).
- A delete function. Prove that delete preserves the binary search tree property and prove a property about the interaction with `elem_of` (*moderate*).
- Self-balancing search trees, such as black-red trees or AVL trees (*very difficult*). See also https://en.wikipedia.org/wiki/Self-balancing_binary_search_tree.

4 Arithmetic expression decompiler

The goal of this project is to formalize a compiler and decompiler for an arithmetical expression language. You have to prove that the compiler is correct and that the decompiler is an inverse of the compiler.

We consider an expression language of arithmetical expressions that is defined by the following inductive types in Coq:

```
Inductive unop :=
  | OSucc : unop
  | OPred : unop.

Inductive binop :=
  | OPlus : binop
  | OMult : binop.

Inductive exp :=
  | ENat : nat -> exp
  | EUnOp : unop -> exp -> exp
  | EBinOp : binop -> exp -> exp -> exp.
```

The expression `ENat` denotes a constant. The unary operator `OSucc` computes the successor and `OPred` computes the predecessor (which is defined to be 0 in case the argument is 0). The operators `OPlus` and `OMult` perform addition and multiplication as usual.

Exercise 4.1 Define a function `eval` that gives a semantics to expressions. Fill out the following template:

```
Fixpoint eval (e : exp) : nat :=
  (* fill out *).
```


Exercise 4.2 Write 2 tests of the `eval` function and prove these tests using the `reflexivity` tactic.

The first part of this project is to define a compiler from arithmetical expressions to a stack machine. The instructions of the stack machine are in reverse Polish notation and are given by the following Coq definitions:

```
Inductive instruction :=
| IPush : nat -> instruction
| IUnOp : unop -> instruction
| IBinOp : binop -> instruction.
```

Exercise 4.3 Implement a compiler:

```
Fixpoint compile (e : exp) : list instruction :=
(* fill out *).
```

Exercise 4.4 Implement a virtual machine for stack instructions.

```
Definition vm : list instruction -> option nat :=
(* fill out *).
```

The virtual machine should yield `None` in case of stack underflow. For example, `vm [EBinOp OPlus; IPush 10] = None`.

Exercise 4.5 Write 2 positive tests and 2 negative tests of the virtual machine and prove these using the `reflexivity` tactic.

Exercise 4.6 Prove correctness of your compiler:

```
Lemma vm_correct : forall e,
  vm (compile e) = Some (eval e).
```

The next part of this project is to define a decompiler. The decompiler translates a list of machine instructions into a corresponding arithmetical expression. You have to show that the decompiler is an inverse of the compiler.

Exercise 4.7 Define a decompiler:

```
Definition decompile : list instruction -> option exp :=
(* fill out *).
```

The decompiler should yield `None` in case the list of instructions is ill-formed. For example, `decompile [EBinOp OPlus; IPush 10] = None`.

The implementation of the decompiler is very similar to the implementation of the virtual machine. However, instead of evaluating a list of instructions to a value, you have to evaluate the list to an expression.

Exercise 4.8 *Prove that the decompiler is a left-inverse of the compiler:*

```
Lemma decompile_correct : forall e,  
  decompile (compile e) = Some e.
```

Exercise 4.9 *Extend your project in one of the following ways:*

- *Extend the language with an operator that may fail (for example, division may fail in case of division by zero) (**moderate**). You have to extend the syntax and semantics of the expression and machine language, as well as your proofs.*
- *Extend the language with a non-deterministic random number generator (**difficult**). You have to extend the syntax and semantics of the expression and machine language, as well as your proofs.*
- *Prove that the decompiler is a right inverse of the compiler (**difficult**):*

```
Lemma compile_decompile : forall ins e,  
  decompile ins = Some e -> compile e = ins.
```

- *As you may have noticed, the decompiler and the virtual machine, as well as its correctness proofs, are very similar. Try to figure out an abstraction so you can factor out similarities (**very difficult**).*