

QuickChick: A Coq Framework For Verified Property Based Testing

Zoe Paraskevopoulou

September 8, 2014

Overview

- **Goal:** Trust high-level logical propositions instead of executable testing code
 - ♦ Gain confidence that the right conjecture is being tested
 - ♦ Gain confidence that the testing is thorough
- **Means:** Automatically map executable testing code (*checkers*) to logical propositions
 - ♦ Reasoning about probabilistic programs (*generators*)
 - ▶ map them to sets of outcomes
- **Evaluation:** Application to a number of sizable case studies
 - ♦ Modularity
 - ♦ Scalability
 - ♦ Minimal changes to existing code

Property-Based Testing

- Popularized by QuickCheck in the FP community
- Achieves a high level of automation
 - ♦ Randomly Generated Data
 - ▶ No need to maintain test suites
 - ♦ The programs are being tested against specifications
 - ▶ No need for human oracle
- The user has to write
 - ♦ Generators
 - ▶ Fine-tuning
 - ▶ Generation of used defined data types
 - ♦ Checkers
 - ▶ Programs that test the desired specification

QuickChick

- Randomized property-based testing framework for Coq
 - ♦ More precisely a port of QuickCheck in Coq
- Checkers need to be executable
 - ♦ No way to directly refute proof goals
- Written in Gallina
- Uses extraction to OCaml for:
 - ♦ Acquiring random seeds
 - ♦ Efficient execution
 - ♦ Generation of numeric (`nat` and `Z`) and boolean values

The value of counterexamples

- QuickChick returns counterexamples for falsifiable conjectures
 - ◆ Support for shrinking
 - ▶ try to isolate the part of the failing input that triggers the failure
- A counterexample could indicate:
 - ◆ errors in program
 - ▶ fix bug
 - ◆ errors in specifications
 - ▶ reformulate checker
- Valuable feedback to understand and fix errors

But ...

How much confidence can we have about the program under test adhering its specifications, when the testing cannot find any more bugs?

Reasons for inadequate testing

Bugs in generators

- May fail to cover sufficiently the input space
 - ♦ Some counterexamples may never generated
- May fail to satisfy preconditions of conditional specifications
 - ♦ A big portion of the generated data can be discarded

Bugs in checkers

- May fail to capture the desired specifications
 - ♦ Too strong preconditions, faulty definitions, . . .

Idea

- Extend QuickChick so it automatically relates checkers to Coq propositions that capture the conjecture under test.
- Manually prove these propositions equivalent to the desired high-level declarative specifications



Guarantee

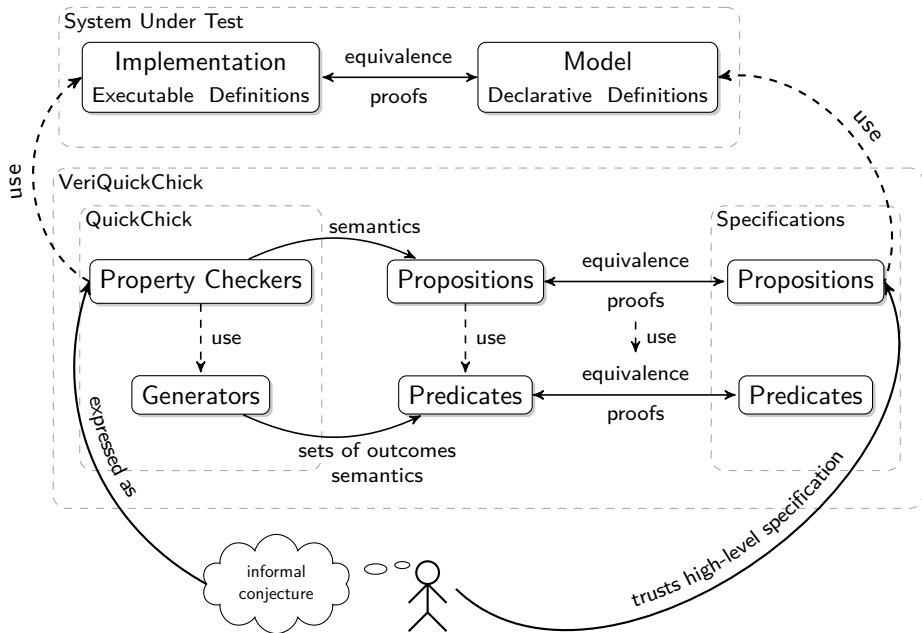
If we could enumerate the output space of the generators used to by the checker without producing any counter examples then we would have a proof by exhaustion for the desired high-level specification

Strategy I

- Map generators to *sets of outcomes*, i.e. the sets of values that have non-zero chance of being generated
 - ♦ Use logical predicates to represent sets
 - ▶ An element $a : A$ belongs to a set that is represented by $P : A \rightarrow \text{Prop}$ if and only if $P\ a$
- Map checkers to logical propositions using the sets of outcomes of the generators that they use
 - ♦ If a generator G for type A is mapped to $S : A \rightarrow \text{Prop}$ then $\text{forAll } G\ f$ is mapped to $\forall x, S\ x \rightarrow f\ x = \text{true}$

Strategy II

- Use the sets of outcomes semantics of generators to prove:
 - ♦ *soundness*
 - ▶ All the values that are generated satisfy a certain predicate
 - ♦ *completeness*
 - ▶ All the values that satisfy a certain predicate can be generated
 - ♦ *correctness*
 - ▶ soundness + completeness
- Use the logical predicates obtained from checkers to prove that they correspond the desired declarative specification



Generators

- Generators are written using a library of combinators (e.g. bind, return, elements, frequency, . . .)
 - ♦ *Primitive* combinators: they depend from the internal generator representation
 - ♦ *Non-primitive* combinators: they are built on top of other combinators
- *Overload* combinators with two kinds of semantics
 - ♦ actual generation semantics
 - ♦ sets of outcomes semantics
- Abstract from the generator representation: Make generators parametric in the generator type constructor
 - ♦ Instantiate them with the set representation to map them to sets of outcomes
 - ♦ or with the actual generator representation to generate data

Set Representation

```
1 Definition Pred (A : Type) : Type := A → Prop.  
2  
3 Definition set_eq {A} (m1 m2 : Pred A) := ∀ A, m1 A ↔ m2 A.  
4  
5 Infix "⟷" := set_eq.
```

- Very compact set representation
 - ♦ Easily models infinite sets
 - ♦ Proof-oriented: facilitates reasoning for set membership
- Primitive combinators need to be implemented differently for each type constructor

Type classes to the rescue

```
1 Class GenMonad (M : Type → Type) :=  
2   {  
3     bindGen : ∀ {A B : Type}, M A → (A → M B) → M B;  
4     returnGen : ∀ {A : Type}, A → M A;  
5     fmapGen : ∀ {A B : Type}, (A → B) → M A → M B;  
6     choose : ∀ {A : Type} {Random A}, A * A → M A;  
7     sized : ∀ {A : Type}, (nat → M A) → M A;  
8     suchThatMaybe : ∀ {A : Type}, M A → (A → bool) →  
9                       M (option A);  
10  }.
```

- All the primitive combinators are included in the type class
- Both generator type constructors are instances of this type class
- The set of outcome definitions of the primitive combinators is **axiomatic**
- The methods of the type class are implicitly parameterized by the type constructor and the corresponding instance

Axioms

$$\text{returnGen } a \equiv \{ x \mid x = a \}$$

$$\text{bindGen } G f \equiv \{ x \mid \exists g, G g \wedge f g x \} \longleftrightarrow \bigcup_{g \in G} f g$$

$$\text{fmapGen } f G \equiv \{ x \mid \exists g, G g \wedge x = f g \}$$

$$\text{choose } (lo, hi) \equiv \{ x \mid lo \leq x \leq hi \}$$

$$\text{sized } f \equiv \{ x \mid \exists n, f n x \} \longleftrightarrow \bigcup_{n \in \mathbb{N}} f n$$

$$\begin{aligned} \text{suchThatMaybe } g P \equiv \{ x \mid x = \text{None} \vee \\ \exists y, x = \text{Some } y \wedge g y \wedge P y \} \end{aligned}$$

Non-primitive combinators

- Non-primitive combinators are built on top of the interface provided by the type class
- We move the definitions of non-primitive combinators inside a section in which we assume in context a type constructor which is instance of the **AbstractGen** type class.
 - ♦ No modification to their code is required
- Primitive combinators used are automatically instantiated with the type constructor and the instance assumed in context

Non-primitive combinators

Example

```
1 Section Utilities.
2   Context {Gen : Type → Type}
3           {H : GenMonad Gen}.
4
5   Definition oneof {A : Type} (def: Gen A) (gs : list (Gen A))
6   : Gen A :=
7     bindGen (choose (0, length gs - 1)) (fun n =>
8       nth def gs n).
9
10    ...
11    ...
12
13 End Utilities.
```

Lemma Library for Non-primitive Combinators

- Using the sets of outcomes semantics we prove correctness for each non-primitive generator combinator
- These lemmas can be used in proofs about user defined generators that use the combinators
 - ♦ less proof duplication and reusability
 - ♦ independence from the implementation of combinators
 - ♦ compositional and more robust proofs

Lemma Library for Non-primitive Combinators

Examples

```
1  Lemma vectorOf_equiv:
2     $\forall \{A : \text{Type}\} (k : \text{nat}) (g : \text{Pred } A),$ 
3       $\text{vectorOf } k \ g \longleftrightarrow \text{fun } l \Rightarrow (\text{length } l = k \wedge \forall x, \text{In } x \ l \rightarrow g \ x).$ 
4
5  Lemma listOf_equiv:
6     $\forall \{A : \text{Type}\} (g : \text{Pred } A),$ 
7       $\text{listOf } g \longleftrightarrow \text{fun } l \Rightarrow (\forall x, \text{In } x \ l \rightarrow g \ x).$ 
8
9  Lemma elements_equiv:
10    $\forall \{A\} (l : \text{list } A) (\text{def} : A),$ 
11      $(\text{elements } \text{def } l) \longleftrightarrow (\text{fun } e \Rightarrow \text{In } e \ l \vee (l = \text{nil} \wedge e = \text{def})).$ 
12
13 Lemma frequency_equiv:
14    $\forall \{A\} (l : \text{list } (\text{nat} * \text{Pred } A)) (\text{def} : \text{Pred } A),$ 
15      $(\text{frequency } \text{def } l) \longleftrightarrow$ 
16        $\text{fun } e \Rightarrow (\exists (n : \text{nat}) (g : \text{Pred } A),$ 
17          $\text{In } (n, g) \ l \wedge g \ e \wedge n <> 0) \vee$ 
18          $((l = \text{nil} \vee \forall x, \text{In } x \ l \rightarrow \text{fst } x = 0) \wedge \text{def } e).$ 
19
20 ...
21 ...
```

Checkers

- Checkers are essentially generators of testing results
- We map them to a proposition that holds iff all the results that belong to the sets of outcomes are successful
 - ♦ The result of a test input is successful if it is equal with the expected
- The simplest form of checkers are boolean predicates
- More complex checkers can be written by utilizing property combinators
 - ♦ change the expected outcome, change default generators, instrumentation
 - ♦ We provide a library of correctness lemmas

Checkers

- Checkers are represented internally with the type operator `Property`.

```
1 Definition Property (Gen : Type → Type) := Gen QProp.
```

- We use the function `semProperty` to map them to logical propositions

```
1 Definition semProperty (P : Property Pred) : Prop :=  
2   ∀ qp, P qp → success qp = true.
```

Testable type class

- **Testable** type class provides a canonical way of turning types that can be tested into a **Property**
- Anything testable can be mapped to a proposition

```
1 Class Testable {Gen : Type → Type} (A : Type) : Type :=  
2   {  
3     property : A → Property Gen  
4   }.  
5  
6 Definition semTestable {A : Type} {_ : Testable A} (a : A)  
7 : Prop :=  
8   semProperty (property a).
```

Lemma Library for Checker Combinators I

forall and implication Lemmas

```
1 Lemma semForAll :
2   ∀ {A prop : Type} {H1 : Testable prop} {H2 : Show A} (gen : Pred A)
3     (f : A → prop),
4     semProperty (forall gen f) ↔ ∀ a : A, gen a → semTestable (f a).
5
6 Lemma semImplication:
7   ∀ {prop : Type} {H : Testable prop} (p : prop) (b : bool),
8     semProperty (b ==> p) ↔ b = true → semTestable p.
```

Lemma Library for Checker Combinators II

Lemmas for specific testable types

```
1 Lemma semBool:
2   ∀ (b : bool), semTestable b ↔ b = true.
3
4 Lemma semFun:
5   ∀ {A prop : Type} {H1 : Show A} {H2 : Arbitrary A} {H3 : Testable prop}
6     (f : A → prop),
7     semTestable f ↔ ∀ (a : A), arbitrary a → semTestable (f a).
8
9 ...
10 ...
```

arbitrary here is a generator for elements of type A. It is a method of the Arbitrary type class that provides a common interface for generation. Testable type class use by default these generators to derive a Property

Lemma Library for Checker Combinators III

Identity Lemmas

Some combinators do not affect the testing outcome. They are used for instrumentation purposes.

```
1 Lemma semCallback_id:
```

```
2   ∀ {prop : Type} {H : Testable prop} (cb : Callback) (p : prop),  
3     semProperty (callback cb p) ↔ semTestable p.  
4
```

```
5 Lemma semWhenFail_id:
```

```
6   ∀ {prop : Type} {H : Testable prop} (s : String.string) (p : prop),  
7     semProperty (whenFail s p) ↔ semTestable p.  
8
```

```
9 Lemma semPrintTestCase_id:
```

```
10  ∀ {prop : Type} {H : Testable prop} (s : String.string) (p : prop),  
11    semProperty (printTestCase s p) ↔ semTestable p.
```

A large, leafless tree silhouette is centered against a solid red background. The tree's branches are intricate and spread out, filling much of the frame. Overlaid on the tree is the text 'Case Study' and 'Red-black Trees' in white.

Case Study

Red-black Trees

Red-Black Trees

- A self-balancing data structure
- Binary trees with an additional color label to each node

```
1 Inductive color := Red | Black.  
2 Inductive tree :=  
3   | Leaf : tree  
4   | Node : color → tree → nat → tree → tree.
```

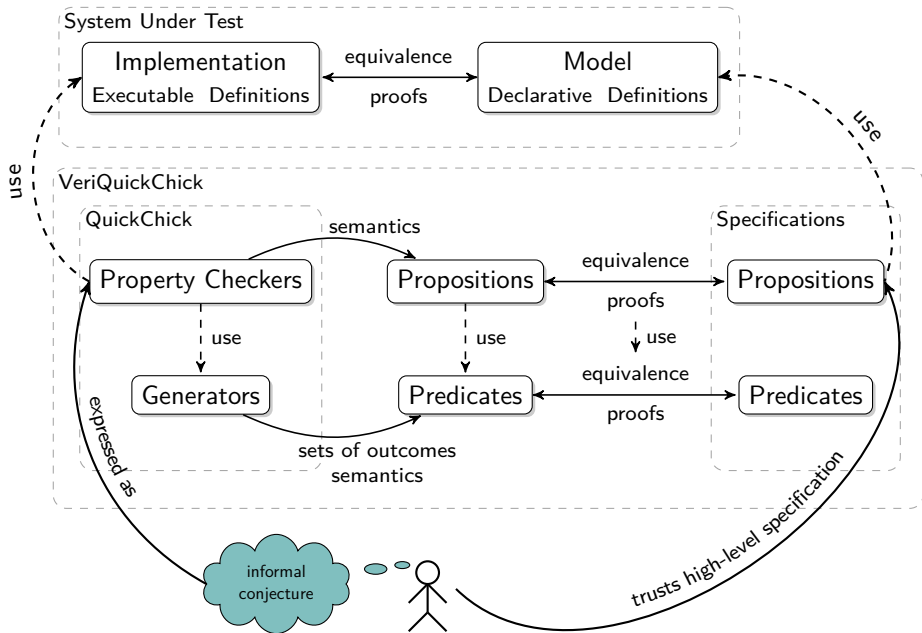
- The should preserve the following invariants
 - ♦ The root is always black
 - ♦ The leaves are empty and black
 - ♦ For each node the path to each possible leaf has the same number of black nodes
 - ♦ Red nodes can only have black children
- If the invariants are preserved then the longest path from the root is at most two times longer than the shortest
- The `insert` method should preserve the invariant
 - ♦ We want to test that!

Red-Black Trees

- A self-balancing data structure
- Binary trees with an additional color label to each node

```
1 Inductive color := Red | Black.  
2 Inductive tree :=  
3   | Leaf : tree  
4   | Node : color → tree → nat → tree → tree.
```

- The should preserve the following invariants
 - ♦ The root is always black
 - ♦ The leaves are empty and black
 - ♦ For each node the path to each possible leaf has the same number of black nodes
 - ♦ Red nodes can only have black children
- If the invariants are preserved then the longest path from the root is at most two times longer than the shortest
- The **insert** method should preserve the invariant
 - ♦ We want to test that!

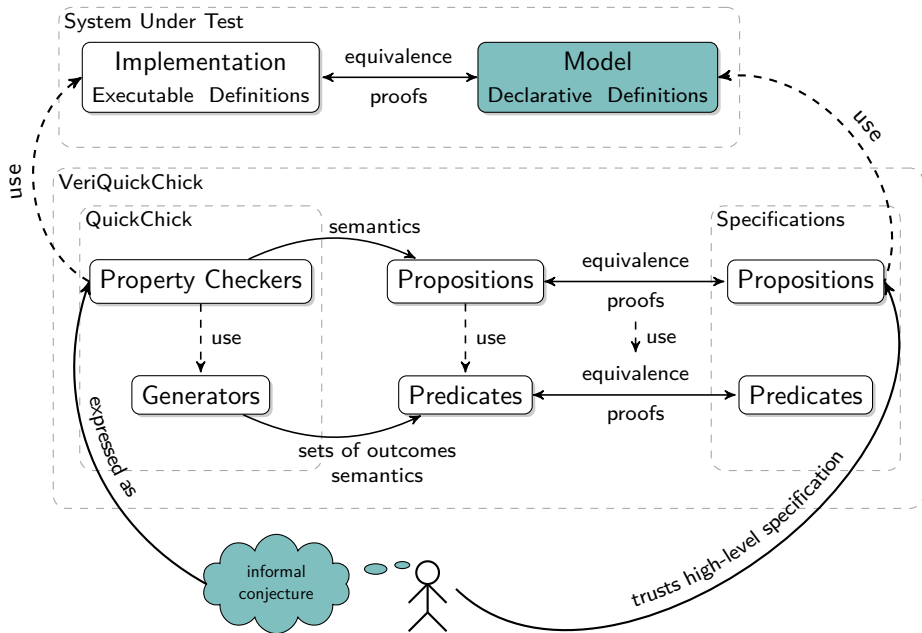


The Red-Black Invariant

```
1 Inductive is_redblack: tree → color → nat → Prop :=  
2   | IsRB_leaf: ∀ c, is_redblack Leaf c 0  
3   | IsRB_r: ∀ n tl tr h,  
4       is_redblack tl Red h → is_redblack tr Red h →  
5       is_redblack (Node Red tl n tr) Black h  
6   | IsRB_b: ∀ c n tl tr h,  
7       is_redblack tl Black h → is_redblack tr Black h →  
8       is_redblack (Node Black tl n tr) c (S h).
```

- $\text{is_redblack } t \ c \ h$ means that t is a subtree of a well-formed RB tree
 - ♦ in *color-context* c (the color of the parent node)
 - ♦ with *black-height* t (# black nodes in each path to a leaf)
- A tree t satisfies the RB invariant iff:

exists h , $\text{is_redblack } t \ h \ \text{Red}$



The Red-Black Invariant

- `insert` should preserve the invariant:

$$\begin{aligned} \forall x t h, \text{is_redblack } h \text{ Red } t \rightarrow \\ \exists h', \text{is_redblack } h' \text{ Red } (\text{insert } x t) \end{aligned}$$

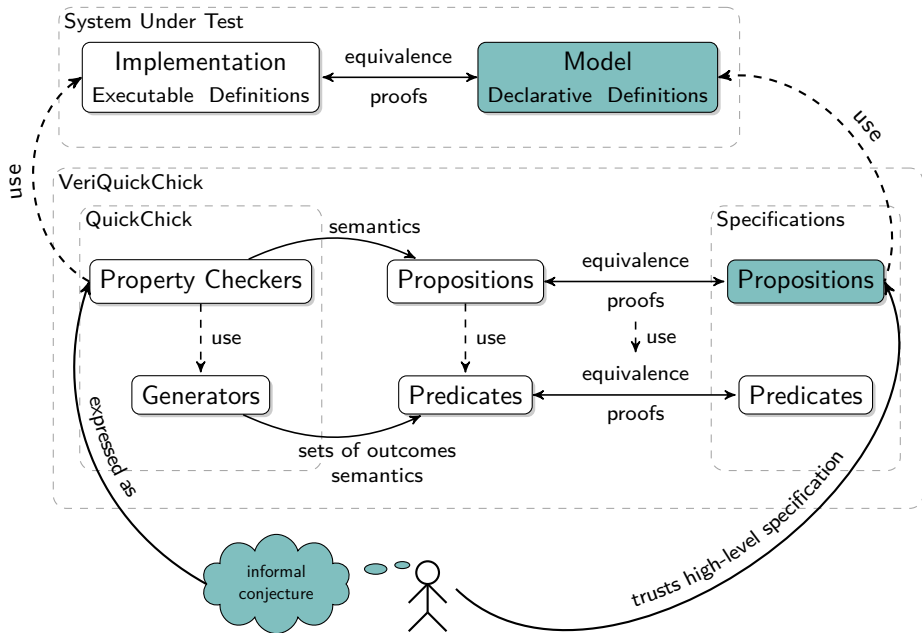
- In order to be able to test this we need
 - ♦ A decision procedure to determine whether a tree satisfies the RB invariant
 - ♦ A generator for RB trees
 - ▶ Should generate only trees that satisfy the invariant
 - ▶ Filtering out ill-formed RB trees is not an option

The Red-Black Invariant

- `insert` should preserve the invariant:

$$\begin{aligned} \forall x t h, \text{is_redblack } h \text{ Red } t \rightarrow \\ \exists h', \text{is_redblack } h' \text{ Red } (\text{insert } x t) \end{aligned}$$

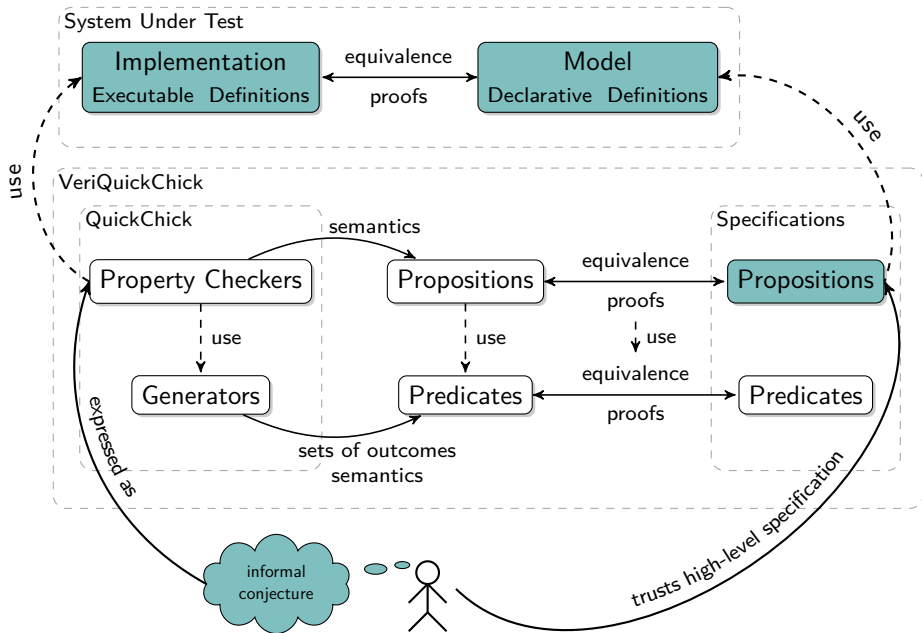
- In order to be able to test this we need
 - ♦ A decision procedure to determine whether a tree satisfies the RB invariant
 - ♦ A generator for RB trees
 - ▶ Should generate only trees that satisfy the invariant
 - ▶ Filtering out ill-formed RB trees is not an option



Executable Definitions

We need a decisional procedure to determine whether a tree satisfies the RB invariant.

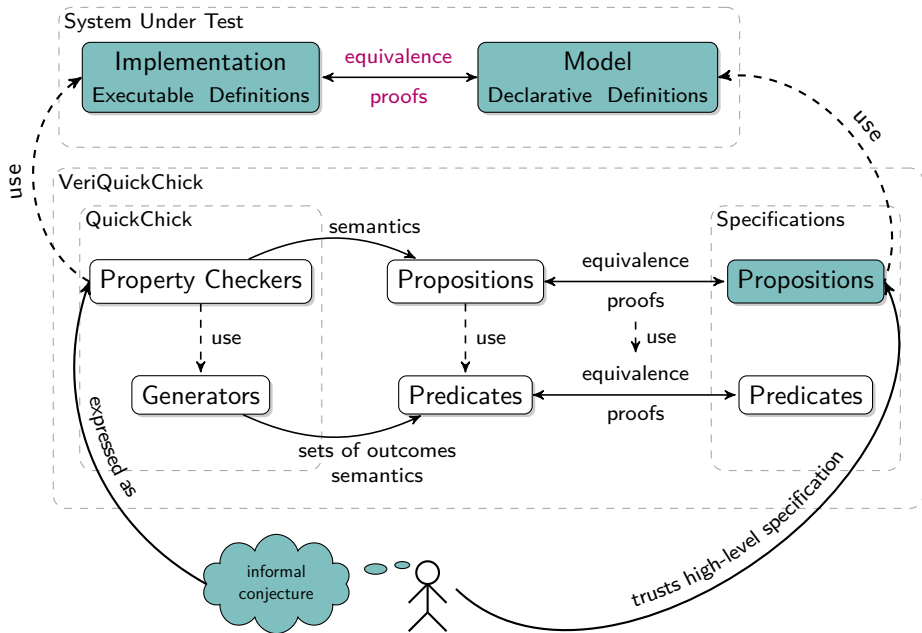
```
1  Fixpoint is_redblack_dec (t : tree) (c: color) : bool :=
2    match t with
3    | Leaf => true
4    | Node c' tl _ tr =>
5      match c' with
6      | Black =>
7        (black_height_dec tl == black_height_dec tr) &&
8        is_redblack_dec tl Black && is_redblack_dec tr Black
9      | Red =>
10       match c with
11       | Black =>
12         (black_height_dec tl == black_height_dec tr) &&
13         is_redblack_dec tl Red && is_redblack_dec tr Red
14       | Red => false
15       end
16     end
17   end.
```



Executable Definitions

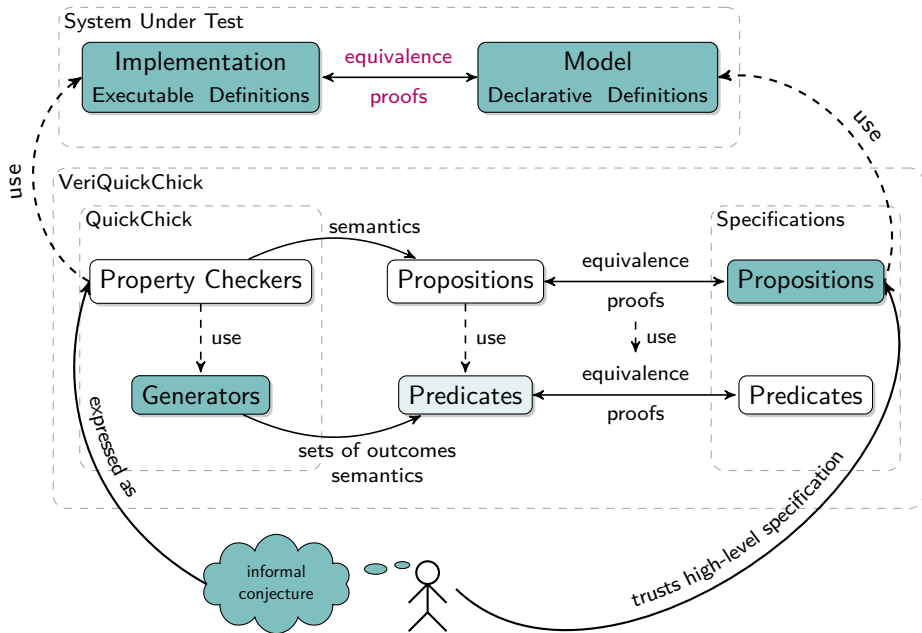
Does it correspond to the inductive definition? Yes!

```
1 Lemma is_redblack_exP :  
2    $\forall$  (t : tree) (c : color),  
3     reflect ( $\exists$  n, is_redblack t c n) (is_redblack_dec t c).
```



Red-black Tree Generator

```
1 Section Generators.
2 Context {Gen : Type → Type}
3         {H: GenMonad Gen}.
4
5 Definition genColor := elements Red [Red; Black].
6
7 Fixpoint genRBTTree_height (h : nat) (c : color) :=
8   match h with
9   | 0 =>
10     match c with
11     | Red => returnGen Leaf
12     | Black => oneof (returnGen Leaf)
13                  [returnGen Leaf;
14                   bindGen arbitraryNat (fun n =>
15                     returnGen (Node Red Leaf n Leaf))]
16     end
17   | S h => ...
18 end.
19
20 Definition genRBTTree := sized (fun h => genRBTTree_height h Red).
21
22 End Generators.
```



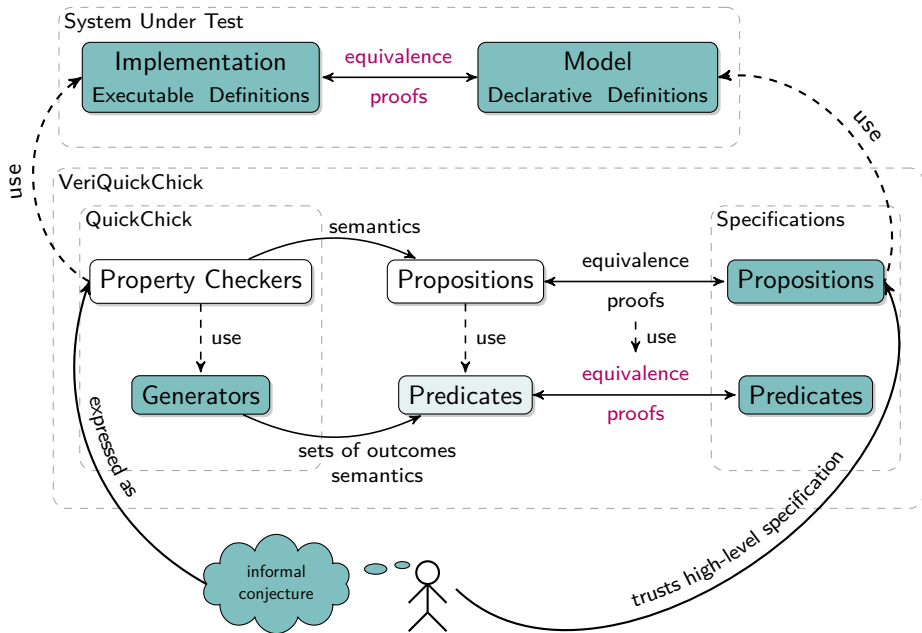
Correctness for Red-black Tree Generator

We want to prove that the generator generates only trees that satisfy the RB invariant and also that it can generate all the possible trees that satisfy the RB invariant.

```
1 Lemma genRBTree_correct:
2   genRBTree  $\longleftrightarrow$  ( $\text{fun } t \Rightarrow \exists h, \text{is\_redblack } t \text{ Red } h$ ).
```

We need an intermediate lemma

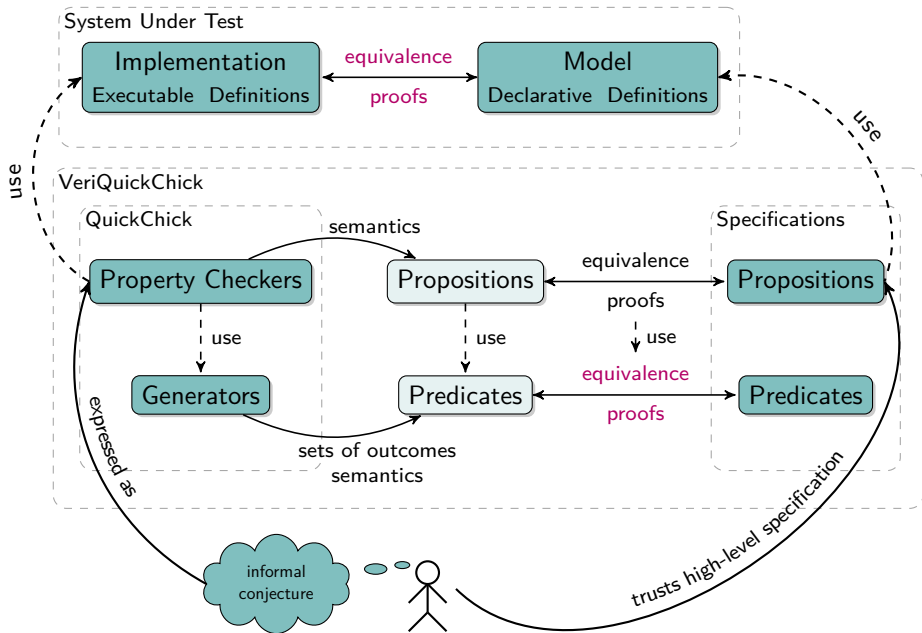
```
1 Lemma genRBTree_height_correct:
2    $\forall c h,$ 
3   ( $\text{genRBTree\_height } h c$ )  $\longleftrightarrow$  ( $\text{fun } t \Rightarrow \text{is\_redblack } t c h$ ).
```



Checker

We can now write the checker for the property:

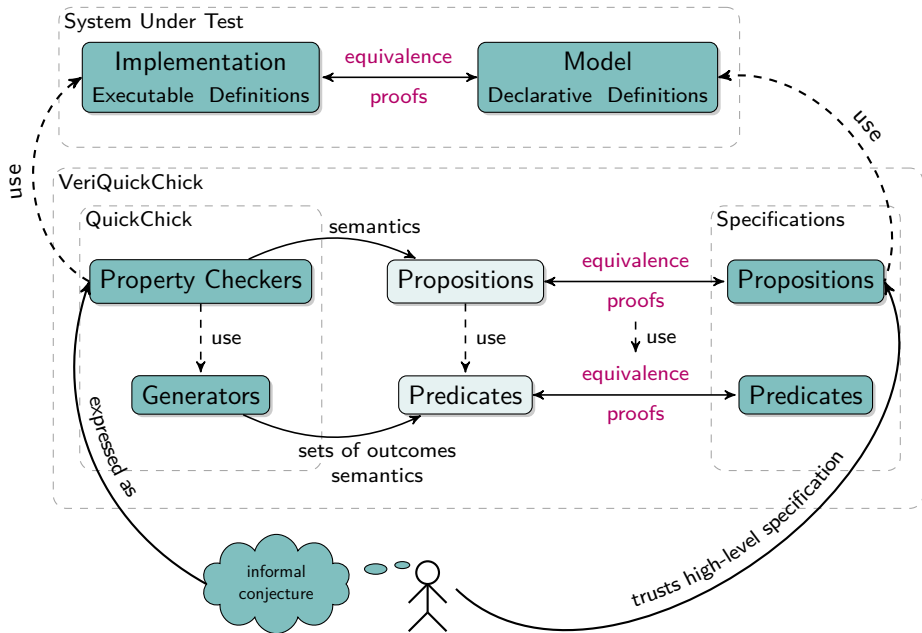
```
1 Section Checker.
2 Context {Gen : Type → Type}
3         {H: GenMonad Gen}.
4
5 Definition insert_is_redblack_checker : Property Gen :=
6   forAll arbitraryNat (fun n =>
7     forAll genRBTree (fun t =>
8       is_redblack_dec t Red ==> is_redblack_dec (insert n t) Red)).
9
10 End Checker.
```



Correctness for Checker

.. and prove that it indeed tests the right thing

```
1 Definition insert_is_redblack :=  
2    $\forall x s h, \text{is\_redblack } s \text{ Red } h \rightarrow \exists h', \text{is\_redblack } (\text{insert } x s) \text{ Red } h'.$   
3  
4 Lemma insert_is_redblack_checker_correct:  
5   semProperty insert_is_redblack_checker  $\leftrightarrow$  insert_is_redblack.
```



Conclusion

- We provide a mechanism to verify that a conjecture under test conforms to a high-level specification
- We facilitate reasoning for probabilistic programs
 - ♦ set of outcomes abstraction
- We proved high-level specifications about combinators
 - ♦ Proofs: 600 LOC
 - ♦ First step towards a fully verified PB testing framework
- We applied our methods to verify complex generators used to test an IFC machine
 - ♦ Generators: 350 LOC / Proofs: 900 LOC
 - ♦ ability to verify existing code
 - ♦ scalability
 - ♦ modularity
- Although reduced, manual effort is still required for the proofs

Future Work

- Remove the axioms by proving that the sets of outcomes of primitive combinators are indeed the those we assume
 - ♦ Fully verified PB testing framework
 - ♦ This would require deeper integration in Coq
 - ▶ Reasoning about random seed and generators of numeric and boolean values
- Verification of a framework for synthesizing generators from specifications
 - ♦ automation + formal guarantees
- Facilitate reasoning for the underlying probability distributions
 - ♦ instantiate generators with probability monad

Thank You!

Questions?

Checkers (internals) I

```
1  Record Result :=
2    MkResult {
3      ok : option bool;  (* Testing outcome *)
4      expect : bool;    (* Expected outcome *)
5      ...               (* Other fields used for tracing *)
6    }.
7
8  Inductive Rose (A : Type) : Type :=
9    MkRose : A → Lazy (list (Rose A)) → Rose A.
10
11 Record QProp : Type := MkProp
12 {
13   unProp : Rose Result
14 }.
15
16 Definition Property (Gen : Type → Type) := Gen QProp.
```

Checkers (internals) II

```
1 Definition resultSuccessful (r : Result) : bool :=
2   match r with
3     | MkResult (Some res) expected _ _ _ =>
4       res == expected
5     | _ => true
6   end.
7
8 Definition success qp :=
9   match qp with
10    | MkProp (MkRose res _) => resultSuccessful res
11  end.
12
13 Definition semProperty (P : Property Pred) : Prop :=
14   ∀ qp, P qp → success qp = true.
```