

Prophecy Variables in Separation Logic

(Extending Iris with Prophecy Variables)

Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy,
Marianna Rapoport, Amin Timany, Derek Dreyer, Bart Jacobs



MAX PLANCK INSTITUTE
FOR SOFTWARE SYSTEMS

What is Iris?! What are prophecy variables?!

The Iris framework:

- ▶ Higher-order concurrent separation logic framework in Coq (developed at MPI-SWS and elsewhere)
- ▶ Deployed in many verification projects (e.g., Rustbelt)
- ▶ Great for verifying tricky concurrent programs

Prophecy variables:

- ▶ Old idea introduced by Abadi and Lamport (1991)
- ▶ Lets you “peek into the future” of a program’s execution when reasoning about an earlier step in the program
- ▶ **Never formally integrated into Hoare logic before!!!**

Our contribution: Prophecy variables for Iris

First integration of prophecy variables to Hoare logic!

- ▶ There was an informal treatment in Vafeiadis's thesis
- ▶ We discovered a flaw in proof of his key example (RDCSS)

Key idea of our approach:

- ▶ Model the right to resolve a prophecy as an ownable resource
- ▶ Leverage separation logic to easily ensure soundness

Implementation in Iris:

- ▶ Everything is formally verified in Coq
- ▶ Including key examples (RDCSS, Herlihy-Wing queues)

Part I – separation logic and prophecy variables

Resources and ownership

The notion of resource is pervasive:

- ▶ File system handle, memory location, permission...
- ▶ A resource should not be used concurrently (but this can sometimes be relaxed)
- ▶ In Iris they are expressed in terms of resource algebras (user-defined structures called cameras)

Examples of resource ownership:

- ▶ $\ell \mapsto v$ (exclusive ownership of a location ℓ)
- ▶ $\ell \overset{q}{\mapsto} v$ (“read only” (fractional) ownership of a location ℓ)
- ▶ $\text{Proph}(p, v)$ (exclusive right to resolve prophecy p)

Separation logic (Iris base logic)

Separating conjunction $P * Q$:

- ▶ The resources are split in disjoint parts satisfying P and Q respectively
- ▶ $l_1 \mapsto v_1 * l_2 \mapsto v_2$ can only be valid if $l_1 \neq l_2$
- ▶ $l \mapsto v_1 * l \mapsto v_2$ is a contradiction
- ▶ The magic wand $P \multimap Q$ is a form of implication

Other logical connectives for building Iris propositions:

- ▶ Conjunction $P \wedge Q$, disjunction $P \vee Q$...
- ▶ Universal and existential quantifiers
- ▶ Ownership of a resource $[a : M]^{\gamma}$ (and related connectives)
- ▶ Persistence modality $\Box P$, later modality $\triangleright P$...

Hoare triples (and weakest preconditions)

We use Hoare triples $\{P\} e \{x.Q\}$ for specifications:

- ▶ The (potential) result of evaluating e is bound in Q
- ▶ Evaluation of e is safe in a state satisfying P
- ▶ If a value is reached the corresponding state and value satisfy Q
- ▶ Defined in terms of weakest preconditions:

$$\{P\} e \{x.Q\} \triangleq \Box(P \rightarrow \text{wp } e \{x.Q\})$$

Weakest preconditions are encoded using the base logic!

Example of specification: eager coin

We consider a specification for a “coin”:

- ▶ A coin is only ever tossed once
- ▶ Reading its value always gives the same result

$$\{\text{True}\} \text{new_coin}() \{x. \exists c. \exists b. x = c \wedge \text{Coin}(c, b)\}$$
$$\{\text{Coin}(c, b)\} \text{read_coin}(c) \{x. x = b \wedge \text{Coin}(c, b)\}$$

Simple (eager) implementation:

$$\text{Coin}(c, b) \triangleq c \mapsto b$$
$$\text{new_coin}() \triangleq \text{ref}(\text{nondet_bool}())$$
$$\text{read_coin}(c) \triangleq !c$$

A lazy coin implementation

What if we want to flip the coin as late as possible?

```
new_coin()  $\triangleq$  ref(None)

read_coin(c)  $\triangleq$  match !c with
  Some(b)  $\Rightarrow$  b
| None     $\Rightarrow$  let b = nondet_bool();
                  c  $\leftarrow$  Some(b); b
end
```

To keep the same spec we need prophecy variables

Specification of the prophecy variables operations

Prophecy variables are used through two ghost code instructions

- ▶ **NewProph** creates a new prophecy variable
- ▶ **Resolve p to v** resolves prophecy variable p to value v

$$\{\text{True}\} \text{NewProph} \{p. \exists v. \text{Proph}(p, v)\}$$
$$\{\text{Proph}(p, v)\} \text{Resolve } p \text{ to } w \{x. x = () \wedge v = w\}$$

Principles of prophecy variables in operation logic:

- ▶ The future is ours
Proph(p, v) gives exclusive right to resolve p
- ▶ We must fulfill our destiny
A prophecy can only be resolved to the predicted value

Back to the lazy coin implementation

```
new_coin()  $\triangleq$  let  $r = \text{ref}(\text{None})$ ;  
               let  $p = \text{NewProph}$ ;  
               { $val = r, \text{proph} = p$ }
```

```
read_coin( $c$ )  $\triangleq$  match ! $c.val$  with  
                 Some( $b$ )  $\Rightarrow b$   
                 | None    $\Rightarrow$  let  $b = \text{nondet\_bool}()$ ;  
                               Resolve  $c.proph$  to  $b$ ;  
                                $c.val \leftarrow \text{Some}(b)$ ;  $b$   
                 end
```

```
Coin( $c, b$ )  $\triangleq$  ( $c.val \mapsto \text{Some } b$ )  
               $\vee$  ( $c.val \mapsto \text{None} * \exists v.$   
                   $\text{Proph}(c.proph, v) * \text{ValToBool}(v) = b$ )
```

Part II – weakest preconditions and adequacy

Model of weakest preconditions in Iris

Encoding of weakest preconditions (simplified):

$$\begin{aligned} \text{wp } e_1 \{ \Phi \} &\triangleq \text{if } e_1 \in \text{Val} \text{ then } \Phi(e_1) \text{ else} && \underline{\text{(return value)}} \\ &\quad \forall \sigma_1. S(\sigma_1) \equiv * \\ &\quad \text{reducible}(e_1, \sigma_1) \wedge && \underline{\text{(progress)}} \\ &\quad \forall e_2, \sigma_2, \vec{e}_f. ((e_1, \sigma_1) \rightarrow (e_2, \sigma_2, \vec{e}_f)) \equiv * \\ &\quad S(\sigma_2) * \text{wp } e_2 \{ \Phi \} * *_{e \in \vec{e}_f} \text{wp } e \{ \text{True} \} && \left. \vphantom{\forall e_2, \sigma_2, \vec{e}_f.} \right\} \underline{\text{(preservation)}} \\ \\ S(\sigma) &\triangleq \boxed{\bullet \sigma}^{\gamma_{\text{heap}}} && \underline{\text{(state interp.)}} \end{aligned}$$

Some intuitions about the involved components:

- ▶ The state interpretation holds the state of the physical heap
- ▶ View shifts $P \equiv * Q$ allow updates to owned resources
- ▶ The actual definition uses the $\triangleright P$ modality to avoid circularity

Operational semantics: head reduction and observations

We extend reduction rules with observations:

$$\begin{aligned}(\bar{n} + \bar{m}, \sigma) &\rightarrow_h (\overline{n + m}, \sigma, \epsilon, \epsilon) \\(\mathbf{ref}(v), \sigma) &\rightarrow_h (l, \sigma \uplus \{l \leftarrow v\}, \epsilon, \epsilon) \\(l \leftarrow w, \sigma \uplus \{l \leftarrow v\}) &\rightarrow_h (l, \sigma \uplus \{l \leftarrow w\}, \epsilon, \epsilon) \\(\mathbf{fork} \{e\}, \sigma) &\rightarrow_h ((), \sigma, e :: \epsilon, \epsilon) \\(\mathbf{Resolve } p \text{ to } v, \sigma) &\rightarrow_h ((), \sigma, \epsilon, (p, v) :: \epsilon) \\(\mathbf{NewProp}, \sigma) &\rightarrow_h (p, \sigma \uplus \{p\}, \epsilon, \epsilon)\end{aligned}$$

A couple of remarks:

- ▶ Observations are only recorded on resolutions
- ▶ The state σ now also records the prophecy variables in scope

Extension for prophecy variables

Encoding of weakest preconditions (simplified):

$$\begin{aligned} \text{wp } e_1 \{ \Phi \} &\triangleq \text{if } e_1 \in \text{Val} \text{ then } \Phi(e_1) \text{ else} && \text{(return value)} \\ &\quad \forall \sigma_1, \vec{\kappa}_1, \vec{\kappa}_2. S(\sigma_1, \vec{\kappa}_1 ++ \vec{\kappa}_2) \equiv * \\ &\quad \text{reducible}(e_1, \sigma_1) \wedge && \text{(progress)} \\ &\quad \forall e_2, \sigma_2, \vec{e}_f. ((e_1, \sigma_1) \rightarrow (e_2, \sigma_2, \vec{e}_f, \vec{\kappa}_1)) \equiv * \\ &\quad S(\sigma_2, \vec{\kappa}_2) * \text{wp } e_2 \{ \Phi \} * *_{e \in \vec{e}_f} \text{wp } e \{ \text{True} \} && \left. \vphantom{S(\sigma_2, \vec{\kappa}_2)} \right\} \text{(preservation)} \end{aligned}$$
$$\begin{aligned} S(\sigma, \vec{\kappa}) &\triangleq [\bullet \sigma.1] \gamma^{\text{heap}} * \exists \Pi. [\bullet \Pi] \gamma^{\text{proph}} \wedge \text{dom}(\Pi) = \sigma.2 \wedge && \text{(state interp.)} \\ &\quad \forall \{ p \leftarrow vs \} \in \Pi. vs = \text{filter}(p, \vec{\kappa}) \end{aligned}$$

Some more intuitions about the involved components:

- ▶ The state interpretation holds observations that remain to be made
- ▶ Observations are removed from the list when taking steps

Statement of safety and adequacy

Safety with respect to a (pure) predicate:

$$Safe_{\phi}(e_1) \triangleq \forall \vec{e}\vec{s}, \sigma, \vec{\kappa}. ([e_1], \emptyset) \rightarrow_{\text{tp}}^* (e_2 :: \vec{e}\vec{s}, \sigma, \vec{\kappa})$$

$$\Rightarrow proper_{\phi}(e_2, \sigma) \wedge \forall e \in \vec{e}\vec{s}. proper_{\text{True}}(e, \sigma)$$

$$proper_{\psi}(e, \sigma) \triangleq (e \in \text{Val} \wedge \psi(e)) \vee \text{reducible}(e, \sigma)$$

Theorem (adequacy). Let e be an expression and ϕ be a (pure) predicate. If $\text{wp } e \{ \phi \}$ is provable then $Safe_{\phi}(e)$.

Conclusion: what I did not show

Our prophecy variables support more features:

- ▶ Multi-resolution prophecies
- ▶ Resolution of prophecies on an atomic instructions
- ▶ Result value included in the prophecy resolutions

Our main motivations for prophecy variables in Iris:

- ▶ Logically atomic specifications (related to linearizability)
- ▶ Allow for much stronger proof rules
- ▶ Examples including RDCSS and Herlihy-Wing queues

Erasure theorem (elimination of ghost code)

Thanks! Questions?

(For more details: <https://iris-project.org>)