

Trillium: Unifying Refinement and Higher-Order Distributed Separation Logic

Léo Stefanescu – *MPI-SWS*

j.w.w. Lars Birkedal, Léon Gondelman, Abel Nieto,
Simon Oddershede Gregersen, and Amin Timany
Aarhus University

A more *intensional* adequacy theorem

“Axiomatic” logics such as Iris are justified by their **adequacy theorems**.

The adequacy for the **standard Hoare triple** essentially guarantees a proved closed program will not crash.

Therefore, **module** specifications are justified by the fact that there are **client program** which are safe.

We want a more **direct** way to argue that a module implements **Paxos**, for example.

Models

We represent the specification as a **labeled transition system** $(\mathcal{M}, \xrightarrow{\ell})$.

Example: **Paxos**, taken from the official TLA+ examples

States: $(\mathcal{S}, \mathcal{B}, \mathcal{V})$ where:

- $\mathcal{S} \in \mathcal{P}(\text{PaxosMessage})$ are the sent messages,
- $\mathcal{B} : \text{Acceptor} \rightarrow \text{BallotNr}_{\perp}$ is the greatest ballot promise;
- $\mathcal{V} : \text{Acceptor} \rightarrow (\text{BallotNr} \times \text{Value})_{\perp}$ is the last accepted value;
- and $\text{BallotNr} := \mathbb{N}$.

Transitions: for example

$$\frac{b \in \text{BallotNr}}{(\mathcal{S}, \mathcal{B}, \mathcal{V}) \rightarrow (\mathcal{S} \cup \{\text{msg1a}(b)\}, \mathcal{B}, \mathcal{V})}$$

Finite traces

A **model trace** of \mathcal{M} is a *non-empty* finite sequence μ of the form

$$\delta_1 \xrightarrow{\ell_1} \delta_2 \xrightarrow{\ell_2} \dots \xrightarrow{\ell_{n-1}} \delta_n$$

We write $\mu \stackrel{\ell}{::} \delta$ for snoc, and define $\text{first}(\mu)$ and $\text{last}(\mu)$.

Program traces have configurations $c = (\text{tp}, \sigma)$ where σ is the **state** and $\text{tp} = [e_1, \dots, e_n]$ is the **thread-pool**. The transitions are lifted from a thread-local reduction relation:

$$\frac{(e_k, \sigma) \rightarrow_h (e'_k, \sigma', \vec{e}_f)}{([e_1, \dots, e_k, \dots, e_n], \sigma) \xrightarrow{\zeta} ([e_1, \dots, e'_k, \dots, e_n, \vec{e}_f], \sigma')}$$

where the **locale** ζ identifies the thread e_k in a language-specific way. We write τ such program traces.

Trillium relates **model traces** and **program traces** with a **simulation**.

The parameters of Trillium

A **model** \mathcal{M} ;

An **Iris predicate** $\text{StateInterp}(\mu, \tau)$ which denote the authoritative ownership of the “current” model and program traces μ and τ ;

A **Coq predicate** $\text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta')$ which defines when the joint evolution

$$\mu \rightsquigarrow \mu \stackrel{\zeta}{\vdash} c' \qquad \tau \rightsquigarrow \tau \stackrel{\ell}{\vdash} \delta'$$

is deemed **valid**.

Example: The Aneris instantiation

Aneris models are not labeled but have an initial state. This defines a map $\text{AnerisModel} \rightarrow \text{Model}$.

The **Aneris state interpretation** is of the form

$$\text{StateInterp}(\mu, \tau) := \text{AnerisStateInterp}(\mu) * \text{OwnAuthModel}(\text{last}(\tau)).$$

Valid transitions are given by the reflexive closure of the transition relation of the model

$$\text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta') := \text{last}(\tau) = \delta' \vee \text{last}(\tau) \rightarrow \delta'.$$

Trillium's notion of refinement

Given a **relation** $\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$, we define

$$\text{Ref}_\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$$

as the greatest-fixpoint of:

$$\text{Ref}_\xi(\mu, \tau) \Leftrightarrow$$

$$\xi(\mu, \tau) \wedge \forall c', \zeta. \text{last}(\mu) \xrightarrow{\zeta} c' \Rightarrow$$

$$(\exists \delta', \ell. \text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta') \wedge \text{Ref}_\xi(\mu \stackrel{\zeta}{::} c', \tau \stackrel{\ell}{::} \delta'))$$

Trillium's notion of refinement

Given a **relation** $\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$, we define

$$\text{Ref}_\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$$

as the greatest-fixpoint of:

$$\text{Ref}_\xi(\mu, \tau) \Leftrightarrow$$

$$\xi(\mu, \tau) \wedge \forall c', \zeta. \text{last}(\mu) \xrightarrow{\zeta} c' \Rightarrow$$

$$(\exists \delta', \ell. \text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta') \wedge \text{Ref}_\xi(\mu \stackrel{\zeta}{::} c', \tau \stackrel{\ell}{::} \delta'))$$

When $\text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta') = \text{last}(\tau) \xrightarrow{\zeta} \delta'$ and $\xi(\mu, \tau) \Leftrightarrow \bar{\xi}(\text{last}(\mu), \text{last}(\tau))$ for some $\bar{\xi}$,

Ref_ξ is the **largest simulation included in** $\bar{\xi}$.

The adequacy theorem

Let e be a **program**, σ a **program state**, and Φ an **Iris predicate** on values. Let δ be a **model state** and $\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$.

Suppose the user-defined relation `ValidEvolution` is **ξ -finitary**. If

$\vdash_{\text{Iris}} \text{StateInterp}(\llbracket [e], \sigma \rrbracket, [\delta]) * \text{wp}_{\top} e @ \zeta \{ \Phi \} * \text{AlwaysHolds}(\xi)$
then $\text{Ref}_{\xi}(\llbracket [e], \sigma \rrbracket, [\delta]) \wedge \text{safe}(e, \sigma)$ **holds in the meta-logic**.

The adequacy theorem

Let e be a **program**, σ a **program state**, and Φ an **Iris predicate** on values. Let δ be a **model state** and $\xi \subseteq \text{ProgTrace} \times \text{ModelTrace}$.

Suppose the user-defined relation ValidEvolution is **ξ -finitary**. If

$$\vdash_{\text{Iris}} \text{StateInterp}(\llbracket [e], \sigma \rrbracket, [\delta]) * \text{wp}_{\top} e @ \zeta \{ \Phi \} * \text{AlwaysHolds}(\xi)$$

then $\text{Ref}_{\xi}(\llbracket [e], \sigma \rrbracket, [\delta]) \wedge \text{safe}(e, \sigma)$ **holds in the meta-logic**.

Where $\text{AlwaysHolds}(\xi)$ is the **Iris predicate**

$$\forall \mu, \tau. \left(\begin{array}{l} \text{first}(\mu) = \llbracket [e], \sigma \rrbracket * \text{first}(\tau) = \delta * \\ \left(\forall \mu', \tau', c', \delta', \zeta, \ell. \mu = \mu' \stackrel{\zeta}{\vdash} c' \wedge \tau = \tau' \stackrel{\ell}{\vdash} \delta' \Rightarrow \right) * \\ \left(\text{ValidEvolution}(\mu', \tau', \zeta, \ell, c', \delta') \wedge \xi(\mu', \tau') \right) \\ \text{StateInterp}(\mu, \tau) * \\ \left(\forall e_1, \dots, e_n, \sigma. \text{last}(\mu) = (e_1, \dots, e_n; \sigma) \Rightarrow \right) \\ \left(\forall 1 \leq i \leq n. e_i \in \text{Val} \vee \text{reducible}(e_i, \sigma) \right) \wedge \\ \left(e_1 \in \text{Val} \Rightarrow \Phi(e_1) \right) \end{array} \right) \stackrel{\top}{\equiv}^* \emptyset \quad \xi(\mu, \tau)$$

The weakest precondition

The Iris predicate $\text{wp}_{\mathcal{E}} e @_{\zeta} \{\Phi\}$ is defined as the **guarded fixpoint**:

$$\left\{ \begin{array}{ll}
\Rightarrow_{\mathcal{E}} \Phi(e) & \text{if } e \in \text{Val} \\
\forall \mu, \tau, K. & \\
\text{thread}(\zeta, \text{last}(\mu)) = K[e] \multimap^* & \\
\text{ValidExec}(\mu) \multimap^* & \\
\text{StateInterp}(\mu, \tau) \stackrel{\mathcal{E}}{\equiv}^* \emptyset & \\
\text{reducible}(e, \text{state}(\text{last}(\mu))) \ast & \\
\forall e', \sigma', \vec{e}_f. & \text{otherwise} \\
(e, \text{state}(\text{last}(\mu))) \rightarrow (e', \sigma', \vec{e}_f) \multimap^* \triangleright \emptyset \stackrel{\mathcal{E}}{\equiv}^* & \\
\exists \delta', \ell. \text{StateInterp}(\mu \dot{\vdash} \text{update}(i, K[e'], \sigma', \text{last}(\mu)), \tau \dot{\vdash} \delta') \ast & \\
\text{ValidEvolution}(\mu, \tau, \zeta, \ell, \text{update}(i, K[e'], \sigma', \text{last}(\mu)), \delta) \ast & \\
\text{wp}_{\mathcal{E}} e' @_{\zeta} \{\Phi\} \ast \bigstar_{1 \leq j \leq k} \text{wp}_{\top} e_{f_j} @_{\zeta_j} \{\text{forkpost}\} &
\end{array} \right.$$

Finiteness requirements

In proving the adequacy theorem, we need to prove **Coq existentials** from the corresponding **Iris existentials** over δ', ℓ .

The semantics of Iris existentials is basically

$$(a, i) \in \llbracket \exists x : A, P \rrbracket \iff \exists v_i \in \llbracket A \rrbracket, (a, i) \in \llbracket P[x := v_i] \rrbracket$$

If the v_i must belong to some **finite set**, by the **pigeon hole principle**,

$$\begin{aligned} a \vDash \exists x : A, P &\iff \forall i, (a, i) \in \llbracket \exists x : A, P \rrbracket \\ &\Rightarrow \exists v \in \llbracket A \rrbracket, a \vDash P[x := v] \end{aligned}$$

Finiteness requirements

In proving the adequacy theorem, we need to prove **Coq existentials** from the corresponding **Iris existentials** over δ', ℓ .

The semantics of Iris existentials is basically

$$(a, i) \in \llbracket \exists x : A, P \rrbracket \iff \exists v_i \in \llbracket A \rrbracket, (a, i) \in \llbracket P[x := v_i] \rrbracket$$

If the v_i must belong to some **finite set**, by the **pigeon hole principle**,

$$\begin{aligned} a \vDash \exists x : A, P &\iff \forall i, (a, i) \in \llbracket \exists x : A, P \rrbracket \\ &\Rightarrow \exists v \in \llbracket A \rrbracket, a \vDash P[x := v] \end{aligned}$$

ValidEvolution is ξ -**finitary** if the following set is finite:

$$\{(\delta', \ell) \mid \text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \ell) \wedge \xi(\mu \stackrel{\zeta}{::} c', \tau \stackrel{\ell}{::} \delta')\}$$

for every μ, τ, ζ, c' .

Example: Paxos

Since Paxos is a distributed algorithm, our aim is to show that the model and the program **send the same messages**:

$$\xi(\mu, \tau) := \exists \mathcal{S}, \text{last}(\mu) = (\mathcal{S}, _, _) \wedge \text{Messages}(\text{last}(\tau)) \sim \mathcal{S}$$

Note that ξ -finiteness is automatic.

We are able to transport the correctness of the model to any execution of the module:

Let e be a distributed system obtained by composing n proposers, m acceptors, and k learners.

For any thread-pool tp and state σ , if $(e, \emptyset) \rightarrow^* (\text{tp}, \sigma)$ and both $\text{Chosen}(\text{Messages}(\sigma), v_1)$ and $\text{Chosen}(\text{Messages}(\sigma), v_2)$ hold then

$$v_1 = v_2.$$

Reasoning rules

The **inference rule** to update the **model state** in an instantiation such as Aneris is:

$$\frac{\delta \rightarrow \delta' \quad \text{Atomic}(e) \quad e \notin \text{Val}}{\text{OwnFragModel}(\delta) * \text{wp}_{\mathcal{E}} e \{x. \text{OwnFragModel}(\delta') * \Phi(x)\} \vdash \text{wp}_{\mathcal{E}} e \{x. \Phi(x)\}}$$

If there exists a certain label $\ell_{stutter}$ such that

$$\text{ValidEvolution}(\mu, \tau, \zeta, \ell_{stutter}, c', \text{last}(\tau))$$

always holds, then **all the inference rules** of the usual Iris weakest precondition hold in Trillium.

Other examples

Hanoi towers and **Incrementing loop**: enforcing memory to go through a succession of values. Uses **events** to detect memory allocations.

Two-phase commit: very similar to the Paxos example

Eventual consistency of CRDTs: proves that if, eventually, no operations are performed, the state of every nodes converge. Makes assumptions on the network and well as the program.

Fair termination of concurrent programs: the rest of this presentation

Fair termination

Termination of **every** execution is too strong a notion for most **concurrent** programs.

Running example:

```
let rec yes b n = if cas b 1 0 then n := !n-1;  
  if n > 0 then yes b n
```

```
let rec no b m = if cas b 0 1 then m := !m-1;  
  if m > 0 then no b m
```

```
let start k = let b = ref 0 in  
  (yes b (ref k) || no b (ref k))
```

This program terminates when the **scheduler** is **fair**.

Fairness

A program trace μ is **fair** if it is finite, or if it is infinite and every **reducible** thread **eventually** takes a step.

A program e is **fairly terminating** if all its **fair traces** are **finite**.

The goal is to use **Trillium** to prove **fair termination** of programs by constructing a **fairness-preserving** and **termination-preserving** refinement with an **abstract model**:



Fairness model

To have a notion of **fairness** for model traces, the models need a surrogate for the notion of thread: **roles**.


A **fairness model** \mathcal{F} is:

1. a set \mathcal{F} of **states**;
2. a set R of **roles**;
3. a labeled **transition** relation $\rightarrow \subseteq \mathcal{F} \times R \times \mathcal{F}$;
4. a finite set $\text{liveroles}(s)$ of **live** roles for each state $s \in \mathcal{F}$;
5. a **fuel bound** $\text{fuelmax}(s) \in \mathbb{N}$ for each state;
6. satisfying certain conditions.

A trace of a **fairness model** is **fair** if it is finite, or if it is infinite and every **role** which is live at some point, **eventually** it takes a step or stops being live.

The Live model construction (1)

The relation \simeq is defined as the relational composition of:



where

- μ is a trace of the program;
- $\bar{\tau}$ is a trace of \mathcal{F} ;
- τ is a trace of a **Trillium model** $\text{Live}(\mathcal{F})$;
- \simeq_f is induced by $\text{Ref}_{\xi_{\text{fair}}}$;
- \simeq_s corresponds to removing **finite stutter**.

The Live model construction (2)

Given a **fairness model** \mathcal{F} , we define a model $\text{Live}(\mathcal{F})$.

Its **states** are triples (s, F, T) of

- a state $s \in \mathcal{F}$;
- a map $F : \text{liveroles}(s) \rightarrow \mathbb{N}$ associating each live role with its **fuel**;
- a map $T : \text{liveroles}(s) \rightarrow \text{Locale}$ which associates each live role with a **locale**.

Its **labels** are $\text{Step}(\rho, \zeta) \mid \text{Stutter}(\zeta)$, where $\zeta \in \text{Locale}$ and ρ is a **role**.

The idea behind the **transitions** is that each thread must share its resources fairly between the **roles** it handles. Each step the thread takes without a step in one of its roles decreases this role's **fuel**.

The Live model construction (3)

There are two kinds of **transitions**:

$$(s, F, T) \xrightarrow{\text{Step}(\rho, \zeta)} (s', F', T')$$

with the (slightly simplified) conditions:

1. $s \xrightarrow{\rho} s'$ in \mathcal{F} ;
2. $T(\rho) = \zeta$;
3. $F'(\rho) \leq \text{fuelmax}(s')$;
4. $\forall \rho' \in R^{-1}(\zeta) \setminus \{\rho\}, F'(\rho') < F(\rho')$.

$$(s, F, T) \xrightarrow{\text{Stutter}(\zeta)} (s, F', T')$$

when $\forall \rho' \in R^{-1}(\zeta), F'(\rho') < F(\rho')$.

Trillium instantiation

$\text{ValidEvolution}(\mu, \tau, \zeta, \ell, c', \delta') := \text{localeof}(\ell) = \zeta \wedge \text{last}(\tau) \xrightarrow{\ell} \delta'$.

The three components of the states of $\text{Live}(\mathcal{F})$ have corresponding **Iris resources**:

- $\text{Modells}(s)$ for the underlying state of \mathcal{F} ;
- $\text{FuelsAre}(\zeta, fs)$ where fs is a finite partial map from **roles** to **fuels**.


Theorem Given a program e , a finitely branching fairness model \mathcal{F} , a state $\delta_0 \in \mathcal{F}$, if

$\text{Modells}(s_0) \multimap \text{FuelsAre}(\zeta_0, fs_0) \multimap \top \Vdash \top \text{wp}_{\top} e @ \zeta_0 \{ \text{FuelsAre}(\zeta_0, \emptyset) \}$

holds in Iris for any fs_0 , and if \mathcal{F} is fairly terminating, then e is fairly terminating.

Back to the example

We relate it with the following **fairness model**:



It is **fairly terminating** because we there is a **well-founded** order such that, for each state, there exists a **role** which makes this order decrease.

Therefore, the program `start` is fairly terminating.

Future work

Compositionality of the models: lots of techniques we could adopt.

Liveness of distributed systems: adapt the logic for fair termination of concurrent programs to prove liveness properties of distributed systems.

Future work

Compositionality of the models: lots of techniques we could adopt.

Liveness of distributed systems: adapt the logic for fair termination of concurrent programs to prove liveness properties of distributed systems.

Thank you!
Questions?