

Time credits and time receipts in Iris

Glen Mével, Jacques-Henri Jourdan, François Pottier

Inria
CNRS, LRI, Univ. Paris Sud, Université Paris-Saclay

April 8, 2019
Prague

This talk

- recent works: **time credits**
aim: prove an upper bound on the running time of a program
- this talk: **time receipts**
aim: **assume** an upper bound on the running time of a program

These are dual notions.

This talk

- recent works: **time credits**
aim: prove an upper bound on the running time of a program
- this talk: **time receipts**
aim: **assume** an upper bound on the running time of a program

These are dual notions.

Example: a unique symbol generator

The function *genSym* returns fresh symbols:

```
let lastSym = ref 0

let genSym() =
  lastSym := !lastSym + 1;
  !lastSym
```

Example: a unique symbol generator

The function *genSym* returns fresh symbols:

```
let lastSym = ref 0 (* unsigned 64-bit integer *)

let genSym() =
  lastSym := !lastSym + 1; (* may overflow! *)
  !lastSym
```

Strictly speaking, this code is **not** correct.

Example: a unique symbol generator

The function *genSym* returns fresh symbols:

```
let lastSym = ref 0 (* unsigned 64-bit integer *)

let genSym() =
  lastSym := !lastSym + 1; (* may overflow! *)
  !lastSym
```

Strictly speaking, this code is **not** correct.

We still want to prove that this code is “correct” in some sense.

The Bounded Time Hypothesis [Clochard **et al.**, 2015]

Counting from 0 to 2^{64} takes **centuries** with a modern processor.

Therefore, this overflow won't happen in a lifetime.

How to express this informal argument in separation logic?

The Bounded Time Hypothesis [Clochard et al., 2015]

Counting from 0 to 2^{64} takes **centuries** with a modern processor.

Therefore, this overflow won't happen in a lifetime.

How to express this informal argument in separation logic?

In this talk:

- We answer this question using **time receipts**.
- We prove that Iris, extended with time receipts, is **sound**.

A closer look at the problem

Specification of genSym

A specification (in separation logic):

$$P \emptyset * \forall S. \left(\begin{array}{l} \{P \ S\} \\ \text{genSym}() \\ \{\lambda n. n \notin S * P(S \cup \{n\})\} \end{array} \right)$$

for some proposition $P \ S$ which represents:

- the ownership of the generator;
- the fact that S is the set of all symbols returned so far.

Tentative proof of genSym

```
let lastSym = ref 0
```

```
let genSym() =
```

```
  lastSym := !lastSym + 1;
```

```
  !lastSym
```

Tentative proof of genSym

 $\{\}$ `let lastSym = ref 0` $\{P \emptyset\}$ $\{P S\}$ `let genSym() =``lastSym := !lastSym + 1;``!lastSym` $\{\lambda n. n \notin S * P(S \cup \{n\})\}$

Tentative proof of genSym

Invariant: $P S \triangleq \text{lastSym} \mapsto \max S$

{

let lastSym = ref 0

{P ∅}

{P S}

let genSym() =

lastSym := !lastSym + 1;

!lastSym

{λn. n ∉ S * P(S ∪ {n})}

Tentative proof of genSym

Invariant: $P S \triangleq lastSym \mapsto \max S$

{}

let $lastSym = \text{ref } 0$

{ $lastSym \mapsto 0$ }

{ $P \emptyset$ }

{ $P S$ }

let $genSym() =$

{ $lastSym \mapsto \max S$ }

$lastSym := !lastSym + 1;$

{ $lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}}$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}}$ }

$!lastSym$

{ $\lambda n. n \notin S * lastSym \mapsto n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Tentative proof of genSym

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S$

{}

let lastSym = ref 0

{lastSym \mapsto 0}

{P \emptyset }

{P S}

let genSym() =

{lastSym \mapsto max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{264}$ }

Wrong

~~{ $\lfloor \max S + 1 \rfloor_{264} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{264}$ }~~

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

An unpleasant workaround: patch the specification

We may add a precondition to exclude any chance of overflow:

$$P \emptyset * \forall S. \left(\begin{array}{l} \{P S * |S| < 2^{64} - 1\} \\ \text{genSym}() \\ \{\lambda n. n \notin S * P(S \cup \{n\})\} \end{array} \right)$$

This pollutes user proofs with cumbersome proof obligations...
which may even be unprovable!

Time receipts in action

Time receipts in separation logic

To count execution steps, we introduce **time receipts**.

Each step produces **one** time receipt, **and only one**:

$$\{\text{True}\} \quad x + y \quad \{\lambda z. z = \lfloor x + y \rfloor_{2^{64}} * \mathbb{X}1\}_{\mathbb{X}}$$

Time receipts sum up:

$$\underbrace{\mathbb{X}1 * \dots * \mathbb{X}1}_n \equiv \mathbb{X}n$$

But time receipts do **not** duplicate (separation logic):

$$\mathbb{X}1 \not/* \mathbb{X}1 * \mathbb{X}1$$

Therefore, $\mathbb{X}n$ is a witness that (at least) n steps have been taken.

Proof of genSym using time receipts

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S$

{}

let lastSym = ref 0

{lastSym \mapsto 0}

{P \emptyset }

{P S}

let genSym() =

{lastSym \mapsto max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}}$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}}$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts

Invariant: $P S \triangleq lastSym \mapsto \max S * \mathbb{X}(\max S)$

{}

let $lastSym = \text{ref } 0$

{ $lastSym \mapsto 0$ }

{ $P \emptyset$ }

We keep track of elapsed time.

{ $P S$ }

let $genSym() =$

{ $lastSym \mapsto \max S$ }

$lastSym := !lastSym + 1;$

{ $lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}}$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}}$ }

! $lastSym$

{ $\lambda n. n \notin S * lastSym \mapsto n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S * \boxtimes(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * \boxtimes 0}

{P \emptyset }

{P S}

let genSym() =

{lastSym \mapsto max S * \boxtimes max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \boxtimes(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \boxtimes(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n * \boxtimes n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts

Invariant: $P S \triangleq lastSym \mapsto \max S * \mathbb{X}(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * $\mathbb{X}0$ }

{P \emptyset }

Initialization

We obtain 0 time receipts for free.

{P S}

let genSym() =

{lastSym \mapsto max S * \mathbb{X} max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \mathbb{X}(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \mathbb{X}(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * lastSym \mapsto n * \mathbb{X}n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S * \mathfrak{X}(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * $\mathfrak{X}0$ }

{P \emptyset }

{P S}

let genSym() =

{lastSym \mapsto max S * $\mathfrak{X}\max S$ }

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \mathfrak{X}(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \mathfrak{X}(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n * \mathfrak{X}n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Tick

Addition produces one time receipt.

The Bounded Time Hypothesis with time receipts

Let N be an arbitrary integer.

We posit the **Bounded Time Hypothesis**:

$$\mathbb{X} N \vdash \text{False}$$

In other words, we assume that no execution lasts for N steps.

The larger N , the weaker this assumption.

Consequence:

$$\mathbb{X} n \vdash n < N$$

Proof of genSym using time receipts and the BTH

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S * \boxtimes(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * \boxtimes 0}

{P \emptyset }

{P S}

let genSym() =

{lastSym \mapsto max S * \boxtimes max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \boxtimes(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \boxtimes(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n * \boxtimes n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts and the BTH

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S * \mathbb{X}(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * \mathbb{X} 0}

{P \emptyset }

Bounded Time

$\mathbb{X}(\max S + 1)$ entails $\max S + 1 < N$.

{P S}

let genSym() =

{lastSym \mapsto max S * \mathbb{X} max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \mathbb{X}(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \mathbb{X}(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n * \mathbb{X}n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts and the BTH

Invariant: $P\ S \triangleq \text{lastSym} \mapsto \max S * \underline{\Sigma}(\max S)$

{}

let lastSym = ref 0

{lastSym \mapsto 0 * $\underline{\Sigma}$ 0}

{P \emptyset }

Bounded Time

We further require $N \leq 2^{64}$.

{P S}

let genSym() =

{lastSym \mapsto max S * $\underline{\Sigma}$ max S}

lastSym := !lastSym + 1;

{lastSym \mapsto $\lfloor \max S + 1 \rfloor_{2^{64}} * \underline{\Sigma}(\max S + 1)$ }

{ $\lfloor \max S + 1 \rfloor_{2^{64}} \notin S * \text{lastSym} \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \underline{\Sigma}(\max S + 1)$ }

!lastSym

{ $\lambda n. n \notin S * \text{lastSym} \mapsto n * \underline{\Sigma}n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }

Proof of genSym using time receipts and the BTH

Invariant: $P S \triangleq lastSym \mapsto \max S * \Sigma(\max S)$

{}

let $lastSym = \text{ref } 0$

$\{lastSym \mapsto 0 * \Sigma 0\}$

$\{P \emptyset\}$

No overflow

Then, $\max S + 1 < 2^{64}$.

$\{P S\}$

let $genSym() =$

$\{lastSym \mapsto \max S * \Sigma \max S\}$

$lastSym := !lastSym + 1;$

$\{lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \Sigma(\max S + 1)\}$

$\{\max S + 1 \notin S * lastSym \mapsto \max S + 1 * \Sigma(\max S + 1)\}$

$!lastSym$

$\{\lambda n. n \notin S * lastSym \mapsto n * \Sigma n\}$

$\{\lambda n. n \notin S * P(S \cup \{n\})\}$

Proof of genSym using time receipts and the BTH

Invariant: $P S \triangleq lastSym \mapsto \max S * \mathbb{X}(\max S)$

{}

let $lastSym = \text{ref } 0$

{ $lastSym \mapsto 0 * \mathbb{X}0$ }

{ $P \emptyset$ }

{ $P S$ }

let $genSym() =$

{ $lastSym \mapsto \max S * \mathbb{X} \max S$ }

$lastSym := !lastSym + 1;$

{ $lastSym \mapsto \lfloor \max S + 1 \rfloor_{2^{64}} * \mathbb{X}(\max S + 1)$ }

{ $\max S + 1 \notin S * lastSym \mapsto \max S + 1 * \mathbb{X}(\max S + 1)$ }

! $lastSym$

{ $\lambda n. n \notin S * lastSym \mapsto n * \mathbb{X} n$ }

{ $\lambda n. n \notin S * P(S \cup \{n\})$ }



Iris \mathfrak{X} , a program logic with time receipts

Time receipts satisfy the **Bounded Time Hypothesis**:

$$\mathfrak{X} N \vdash \text{False}$$

Each step produces one time receipt; for instance:

$$\{\text{True}\} \quad x + y \quad \{\lambda z. z = \lfloor x + y \rfloor_{2^{64}} * \mathfrak{X}1\}_{\mathfrak{X}}$$

Iris[⌚], a program logic with time receipts

Time receipts satisfy the **Bounded Time Hypothesis**:

$$\mathbb{N} \vdash \text{False}$$

Each step produces one time receipt; for instance:

$$\{\text{True}\} \quad x + y \quad \{\lambda z. z = \lfloor x + y \rfloor_{2^{64}} * \mathbb{N}1\}_{\mathbb{N}}$$

We can obtain zero time receipts unconditionally:

$$\vdash \mathbb{N}0$$

Time receipts are additive:

$$\mathbb{N}m * \mathbb{N}n \equiv \mathbb{N}(m + n)$$

Soundness of Iris with time receipts

Soundness of Iris_x

We want our program logic Iris_x to satisfy this property:

Theorem (Soundness of Iris_x)

If the following Iris_x triple holds:

$$\{\text{True}\} e \{ _ \}_x$$

then e cannot crash *until N steps have been taken*.

We say that “ e is $(N - 1)$ -safe”.

Crashing means trying to step while in a stuck configuration; for example, dereferencing a non-pointer.

Proof sketch of the soundness theorem

We use Iris as a model of Iris^Σ.

$$\{P\} e \{\varphi\}_{\mathbf{x}} \triangleq \{P\} \langle\langle e \rangle\rangle \{\varphi\}$$

The transformation $\langle\langle \cdot \rangle\rangle$ inserts *ticks* (see next slides).

The proof then works as follows:

$$\{\text{True}\} \langle\langle e \rangle\rangle \{_ \}$$



Soundness theorem of Iris [Jung et al., 2015]

$\langle\langle e \rangle\rangle$ is safe



Simulation lemma

e is $(N - 1)$ -safe

The program transformation

We keep track of the number of steps using a global counter c , initialized with 0.

The transformation inserts one *tick* instruction per operation.

$$\llbracket e_1 + e_2 \rrbracket \triangleq \text{tick} (\llbracket e_1 \rrbracket + \llbracket e_2 \rrbracket)$$

tick increments c . On its N^{th} execution, it does not return.

```
let tick x =
  !c := !c + 1;
  if !c < N then x else loop ()
```

Idea: transform a program that runs for too long into a program that never ends, hence is safe.

The simulation lemma

This program transformation does satisfy the desired lemma:

Lemma (Simulation)

*If $\langle\langle e \rangle\rangle$ is safe (i.e. it cannot crash),
then e is $(N - 1)$ -safe (i.e. it cannot crash until N steps have been taken).*

The model of time receipts

The transformation maintains the invariant $!c < N$.

$\boxtimes 1$ is modeled as an exclusive portion of the value of the counter c (Iris features used: authoritative monoidal resource, invariant).

In particular, $\boxtimes n \vdash !c \geq n$. Hence, $\boxtimes N \vdash \text{False}$.


All other axioms of time receipts are realised as well.

Conclusion

Conclusion

Contributions ([new](#)):

	<i>Soundness</i>	<i>Application</i>
<i>Time credits</i>	✓	Reconstruction of Okasaki and Danielsson's thunks (amortized analysis)
<i>Time receipts (exclusive / persistent)</i>	✓	Reconstruction of Clochard et al.'s overflow-free integers
<i>Time credits and time receipts</i>	✓	Proof of Union-Find: complexity, absence of overflow in ranks

Defined within Iris, machine-checked with Coq 

Open question: Can we prove useful facts about **concurrent** code?

Thank you for your time.



What about concurrency?

Iris is a concurrent separation logic; thus, our program logics already support concurrency: they measure the **work** (total number of operations in all threads).

```
let tick x =  
  if (FAA c 1 < N - 1) then x else loop()
```

What about measuring the **span** (running time of the longest-living thread)?

A path to explore: a separate notion of time receipt for each thread, with a rule to clone time receipts of the calling thread when forking.

Compiling code analysed with time receipts

For time receipt proofs to be valid, we need to forbid optimizations!
Otherwise, programs may compute faster than expected.

For example:

```
for i from 1 to N do
  ()
done;
(* This point is beyond the scope of Iris2:
 * anything below may be unsafe,
 * but it shouldn't be reached in a lifetime... *)
crash ()
```

A compiler may optimize it to:

```
(* Too bad! *)
crash ()
```

A solution: insert actual *tick* operations and make them opaque.

Example application: Union-Find

We implement the Union-Find with ranks stored in machine words. While proving the correctness of the algorithm, we also prove

- its complexity (using time credits)
- and the absence of overflows for ranks (using time receipts).

Granted that $x, y \in D$ and $\log_2 \log_2 N < \text{word_size} - 1$, we show the Iris^{\$\Sigma\$} triple:

$$\{ \text{isUF } D \ R \ V * \$(44\alpha(|D|) + 152)\} \\ \text{union } x \ y \\ \{ \lambda z. \text{isUF } D \ R' \ V' * (z = R \ x \vee z = R \ y) \} \$\Sigma$$

Consequences:

- the (amortized) complexity is the inverse Ackermann function;
- if $N = 2^{64}$, then $\text{word_size} \geq 8$ is enough to avoid overflows.

Example: a unique symbol generator (functional version)

Code:

```
let makeGenSym() =  
  let lastSym = ref 0 in (* unsigned 64-bit integer *)  
  fun () →  
    lastSym := !lastSym + 1; (* may overflow *)  
    !lastSym
```

Specification (in higher-order separation logic):

$$\left\{ \begin{array}{l} \{\text{True}\} \\ \text{makeGenSym}() \\ \lambda \text{ genSym}. \exists P. \\ P \emptyset * \forall S. \left(\begin{array}{l} \{P \ S\} \\ \text{genSym}() \\ \{\lambda n. n \notin S * P(S \cup \{n\})\} \end{array} \right) \end{array} \right\}$$

Alternative specification of makeGenSym

Specification (in Iris):

$$\left\{ \begin{array}{l} \{\text{True}\} \\ \text{makeGenSym}() \\ \lambda \text{ genSym}. \exists \gamma. \\ \forall n. \left(\begin{array}{l} \{\text{True}\} \\ \text{genSym}() \\ \{\lambda m. \text{OwnSym}_\gamma(m)\} \end{array} \right) \end{array} \right\}$$

- The ownership of the generator is shared through an invariant.
- $\text{OwnSym}_\gamma(m)$ asserts uniqueness of symbol m :

$$\text{OwnSym}_\gamma(m_1) * \text{OwnSym}_\gamma(m_2) \multimap m_1 \neq m_2$$

A program logic with time credits

Each step consumes one time credit; for instance:

$$\{\$1\} \quad x + y \quad \{\lambda z. z = \lfloor x + y \rfloor_{2^{64}}\}_{\mathbf{x}}$$

We can obtain zero time credits unconditionally:

$$\vdash \$0$$

Time credits are additive:

$$\$m * \$n \equiv \$(m + n)$$

Our program logic $\text{Iris}^{\$}$ satisfies this property:

Theorem (Adequacy of $\text{Iris}^{\$}$)

If the following $\text{Iris}^{\$}$ triple holds:

$$\{ \$n \} e \{ \varphi \} \$$$

then:

- *e cannot crash;*
- *if e computes a value v, then φv holds;*
- ***e computes for at most n steps.***

Our program logic Iris^{Σ} satisfies this property:

Theorem (Adequacy of Iris^{Σ})

If the following Iris^{Σ} triple holds:

$$\{\text{True}\} e \{\varphi\}_{\Sigma}$$

then:

- e cannot crash **until N steps have been taken**;
- if e computes a value v **in less than N steps**, then φv holds.

A program logic with duplicable time receipts

Duplicable time receipts satisfy the **Bounded Time Hypothesis**:

$$\boxtimes N \vdash \text{False}$$

Each step increments a duplicable time receipt; for instance:

$$\{\boxtimes m\} \quad x + y \quad \{\lambda z. z = \lfloor x + y \rfloor_{2^{64}} * \boxtimes(m + 1)\}_{\mathbf{x}}$$

We can obtain zero duplicable time receipts unconditionally:

$$\vdash \boxtimes 0$$

Duplicable time receipts obey maximum:

$$\boxtimes m * \boxtimes n \equiv \boxtimes \max(m, n)$$

Duplicable time receipts **are** duplicable:

$$\boxtimes m \multimap \boxtimes m * \boxtimes m$$

Relation between time receipts and duplicable time receipts:

$$\boxtimes m \vdash \boxtimes m * \boxtimes m$$

Overflow-free integers (summable)

$$\text{IsClock}(v, n) \triangleq 0 \leq n < 2^{64} * v = n * \text{⌈}n$$

- non-duplicable
- supports addition (consumes its operands):

$$\{\text{IsClock}(v_1, n_1) * \text{IsClock}(v_2, n_2)\}$$

$$v_1 + v_2$$

$$\{\lambda w. \text{IsClock}(w, n_1 + n_2)\}$$

no overflow!

Overflow-free integers (incrementable)

$$\text{lsSnapClock}(v, n) \triangleq 0 \leq n < 2^{64} * v = n * \boxtimes n$$

- duplicable
- supports incrementation (does **not** consume its operand):

$\{\text{lsSnapClock}(v, n)\}$

$v + 1$

$\{\lambda w. \text{lsSnapClock}(w, n + 1)\}$

no overflow!

- prgm is a program (source code).
- Pre and $Post$ are logical formulas.

$\{Pre\} \text{prgm} \{Post\}$

Soundness:

"If Pre holds, then prgm won't crash."

(Partial) correctness:

"If Pre holds, then after prgm is run, $Post$ will hold."

Total correctness:

"If Pre holds, then prgm terminates and, after prgm is run, $Post$ will hold."

P is a resource.

$x \mapsto v$ is an exclusive resource, its ownership cannot be shared.

- Standard logic: $P \Rightarrow P \wedge P$
- Separation logic: $P \not\vdash P * P$ (resources are not duplicable)

$P * Q$ are disjoint resources.

$x \mapsto v * x \mapsto v'$ is absurd.

Affine sep. logic: $P * Q \multimap P$ (resources can be thrown away)

Iris is:

- an affine separation logic,
- higher-order,
- full-featured (impredicative invariants, monoidal resources. . .),
- very extensible,
- formalized in Coq. 