

SEPARATION LOGIC



Derek Dreyer

MPI for Software Systems

Cornell/Maryland/Max Planck Summer School

Saarbrücken, August 2017

Separation Logic in One Slide

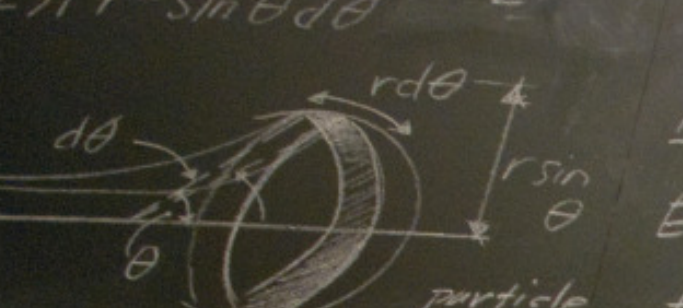
- Extension of **Hoare logic** for reasoning modularly about pointer-manipulating code
- O'Hearn, Reynolds, et al. (~2000)
- One of the most fundamental advances in program verification in the past 20 years
- **“Concurrent separation logic”** (2007) won the **2016 Gödel Prize** (Nobel prize in theory)
- Underpinning of all my recent research

$$1.6 \times 10^{-14} \text{ N}$$

$$N_e \pi n t \left(\frac{1.2 \times 10^2}{8 \pi \epsilon_0 k} \right)^2 \frac{J}{eV}$$

$$2 \pi r^2 \sin \theta d\theta$$

$$\cot \frac{\theta}{2}$$



particle

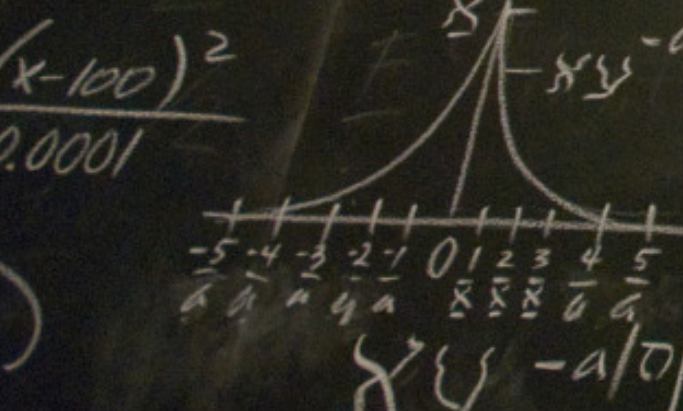
(the point x_j)

TIME INDEPENDENT $V(\vec{r})$

CONSTANT

MOMENTUM \times POSITION UNCERTAINTY GREATER THAN CONSTANT

and $(\hbar = h/2\pi = 6.6 \times 10^{-34})$



$x^2 - a(x)$

$\Delta x = \sqrt{x}$

$\Delta x = \sqrt{x}$

$x^2 - a/\sigma p^2 x$

$$L = \sqrt{l(l+1)} \hbar$$

$$V = \frac{A}{r^n} - \frac{B}{r^m}$$

$$E^{tot}(x_j, t) = \sum_{n \neq j} \frac{E_{ret}^{(n)}(x_j, t) + E_{adv}^{(n)}(x_j, t)}{2}$$

$$E^{free}(x_j, t) = \sum_n E_{ret}^{(n)}$$

$$E^{tot}(x_j, t) = \sum_n \frac{E_{ret}^{(n)}(x_j, t)}{2} + \sum_n \frac{E_{adv}^{(n)}(x_j, t)}{2}$$

$$E^{tot}(x_j, t) = \sum_{n \neq j} E_{ret}^{(n)}(x_j, t) + E_{damping}(x_j, t)$$

$C = \text{speed of light}$

$$\psi(r, t) = \psi(r) e^{-iEt/\hbar}$$

$$\psi(x, t) = \int A(k) e^{i(kx - \omega t)} dk$$

$$\mathcal{E} = \oint_{\partial \Sigma(t)} d\ell \cdot F/q = \oint_{\partial \Sigma(t)} d\ell \cdot (E + v \times B)$$

$$\mathcal{E} = - \frac{d\Phi_B}{dt}$$

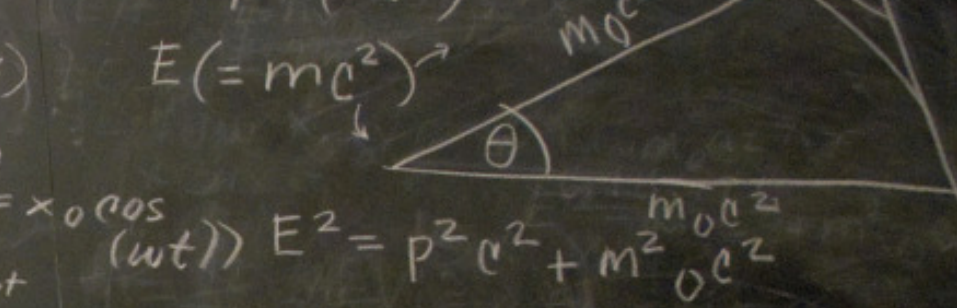
$$F = q(-\nabla \phi - \frac{\partial A}{\partial t})$$

$$E = -\nabla \phi - \frac{\partial A}{\partial t}$$

$$P = \sum_{i=1}^n m_i v_i = m_1 v_1 + m_2 v_2 + m_3 v_3 + \dots + m_n v_n$$

$$(E/c)^2 - p^2 = (m_0 c)^2$$

$$E (= m_0 c^2)$$



$E^2 = p^2 c^2 + m_0^2 c^4$

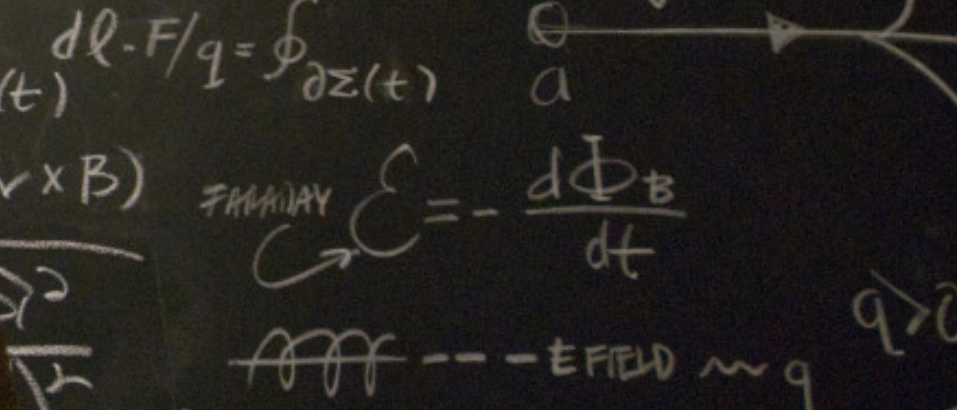
$$P = \frac{h}{\lambda} = \frac{E}{c} \rightarrow E = h\nu \rightarrow E = \hbar \omega$$

$$E_{damping}(x_j, t) = \frac{q}{6\pi\epsilon_0^3} \ddot{x}$$

$$\vec{F} = q[\vec{E} + (\vec{v} \times \vec{B})]$$

$q=0$

$E \cdot B \sim \text{in principle} - \text{finite } E \cdot B$



$$\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$$

$$\vec{F} = q(-\nabla \phi - \frac{\partial A}{\partial t})$$

$$E = -\nabla \phi - \frac{\partial A}{\partial t}$$

RUSTBELT:



LOGICAL FOUNDATIONS FOR THE FUTURE OF SAFE SYSTEMS PROGRAMMING



Derek Dreyer
MPI for Software Systems

*Cornell/Maryland/MPI Summer School
Saarbrücken, August 2017*

Goal of These Lectures

- Tell you about a major ongoing research project in PL/verification (**RustBelt**)
- Teach you something about a cutting-edge language (**Rust**), a cutting-edge separation logic (**Iris**), and how they are connected

About Me

- Born in NYC, grew up in Great Neck
- Undergrad in Math/CS at NYU (1993-1996)
- PhD in CS at CMU (1997-2004)
- Postdoc at TTI-Chicago (2005-2007)
- MPI for Software Systems (2008-present)

About Me

- Started out mainly interested in PL design
 - Particularly “functional” languages (ML, Haskell, etc.)
 - PhD thesis and postdoc work on extensions of the ML module system
- Module systems research was fun, but a bit lonely, and it was hard to have much impact


About Me

- After coming to MPI-SWS, became more interested in foundational PL questions:
 - How can we verify “real” programs?
 - How can we prove safety of “real” PLs?
- Definition of “real” has changed over time...
 - And gotten progressively more “grungy”

**“I AM NOT AN ANIMAL!
I AM A HUMAN BEING!
I...AM...A MAN!”**





A photograph of a living room. In the center, a small, light-colored dog is lying on a dark brown armchair. The chair has an orange and white striped cushion. To the left of the chair is a large potted plant with long, thin green leaves. To the right is a black side table with a potted plant and some items on it. The floor is made of light-colored wood. In the foreground, several small bottles are scattered on the floor. A white text box is overlaid on the image, containing the text "Check out my blog at herrdreyaer.wordpress.com".

Check out my blog at
herrdreyaer.wordpress.com

The RUSTBELT Team

@MPI-SWS & Uds



Ralf
Jung



Jan-Oliver
Kaiser



David
Swasey



Hai
Dang



Jacques-Henri
Jourdan



Ori
Lahav



Viktor
Vafeiadis




Derek
Dreyer

@Mozilla



Aaron
Turon



Niko
Matsakis

@Aarhus



Robbert
Krebbers



Lars
Birkedal

@Seoul Nat. Univ.



Jeehoon
Kang




Chung-Kil
Hur

The RUSTBELT Team

Two new members in July!



Ralf
Jung



Jan
K



Azalea Raad



Josh Yanovski



Derek
Dreyer



Aaron
Turon




Niko
Matsakis



Robbert
Krebbers



Lars
Birkedal



Jeehoon
Kang



Chung-Kil
Hur

@M


Nat. Univ.

Safety vs. Control



Safety vs. Control

Java
C#
Go
Haskell
...



high-level
applications

Safety vs. Control

Java
C#
Go
Haskell
...



high-level
applications



low-level systems
programming

C
C++
...

Safety vs. Control

Java
C#
Go
Haskell
...

high
appl

BOUNTY HUNTERS ATTENTION!
WANTED

A safe
systems
programming
language

\$5,000 REWARD!
NOTIFY NEAREST LAW ENFORCEMENT AGENCY

C
C++
...



systems
ning

Rust:

The Future of Safe Systems Programming?



Rust has been developed at **Mozilla** since 2010

- Mozilla is using Rust to build Servo, a next-gen browser engine with better parallel performance



Rust is the only “systems PL” to provide...

- Low-level control à la modern C++
- Strong safety guarantees
- Industrial development and backing



15 companies using Rust in production

- **Dropbox** is rewriting block storage engine from Go into Rust to control memory footprint

Rust:

The Future of Safe Systems Programming?



Rust has been developed at **Mozilla** since 2010

- Mozilla is using Rust to build Servo, a next-gen browser engine with better parallel performance

**Rust has the potential to become the
“next big thing” in systems programming**




15 companies using Rust in production

- **Dropbox** is rewriting block storage engine from Go into Rust to control memory footprint


Core Idea of Rust




Core Idea of Rust




Core Idea of Rust



Core Idea of Rust



Core Idea of Rust



Unrestricted mutation and aliasing lead to:

- use-after-free errors (dangling references)
- data races
- iterator invalidation

Data Races


$x := 1$

$y := 2$

$x := 3$

Data Races


x := 1
y := 2
x := 3



~~x := 1~~
y := 2
x := 3

Data Races


$x := 1$	\parallel	<code>print y</code>
$y := 2$		<code>print x</code>
$x := 3$		



$x := 1$	\parallel	<code>print y</code>
$y := 2$		<code>print x</code>
$x := 3$		

Data Races


$x := 1$	\parallel	$\text{print } y$
$y := 2$		$\text{print } x$
$x := 3$		



$x := 1$	\parallel	$\text{print } y$
$y := 2$		$\text{print } x$
$x := 3$		$\rightarrow 2 \ 0$

Data Races

$x := 1$	\parallel	$\text{print } y$
$y := 2$		$\text{print } x$
$x := 3$		$\not\rightarrow 2 \ 0$



$x := 1$	\parallel	$\text{print } y$
$y := 2$		$\text{print } x$
$x := 3$		$\rightarrow 2 \ 0$

Data Races

$x := 1$		print y
$y := 2$		print x

Standard compiler optimizations change the
“meaning” of racy concurrent code

$y := 2$		print x
$x := 3$		→ 2 0

Data Races

```
x := 1    ||    print v
```

```
y :
```

Many architectures do, too!

**Standard compiler optimizations change the
“meaning” of racy concurrent code**


```
y := 2    ||    print x  
x := 3    → 2 0
```

What's a PL to do?

Get used to disappointment.

- **Java**: Data races → **Very weak** behavior
- **C/C++**: Data races → **Undefined** behavior


Core Idea of Rust



Unrestricted mutation and aliasing lead to:

- use-after-free errors (dangling references)
- data races
- iterator invalidation

Core Idea of Rust



Rust prevents all these **errors** using a sophisticated “**ownership**” type system

Ownership & Borrowing



- Having a value of type **T** means you “own” it fully.
- **T** can be “borrowed” (e.g. passed by reference):
 - ◆ **&T** — shared, immutable borrow
 - ◆ **&mut T** — unique, mutable borrow

But sometimes you *need*
aliased mutable state!

But sometimes you *need*
aliased mutable state!


Synchronization mechanisms:

- e.g. Locks, channels, semaphores

Memory management:

- e.g. Reference counting

The Reality of Rust



The Reality of Rust

Arc



...standard libraries...

Mutex



```
...  
pub fn borrow(&self) -> Ref<T> {  
    match BorrowRef::new(&self.borrow) {  
        Some(b) => Ref {  
            _value: unsafe { &*self.value.get() },  
            _borrow: b,  
        }, ...  
    }  
}  
...  
...
```

RefCell



Channel



The Reality of Rust

Arc



...standard libraries...

Mutex



Claim of Rust library developers:

Unsafe blocks are safely encapsulated
by their APIs.

RefCell

...

Channel




Is Rust Safe?

Several bugs found in Rust safety so far:

- Due to **unsafe blocks in Rust libraries**
 - e.g. “scoped threads” API
- Due to **dark corners of the type system**
 - e.g. “dropck” rule for checking safety of generic destructor methods



Is Rust Safe?



Pre-Pooping Your Pants With Rust

Alexis Beingessner - April 27, 2015

Leakpocalypse

Much existential anguish and ennui was recently triggered by [Rust Issue #24292: std::thread::JoinGuard \(and scoped\) are unsound because of reference cycles](#). If you feel like you're sufficiently familiar with Leakpocalypse 2k15, feel free to skip to the next section. If you've been thoroughly stalking all my online interactions, then you've basically seen everything in this post already. Feel free to close this tab and return to scanning my IRC logs.

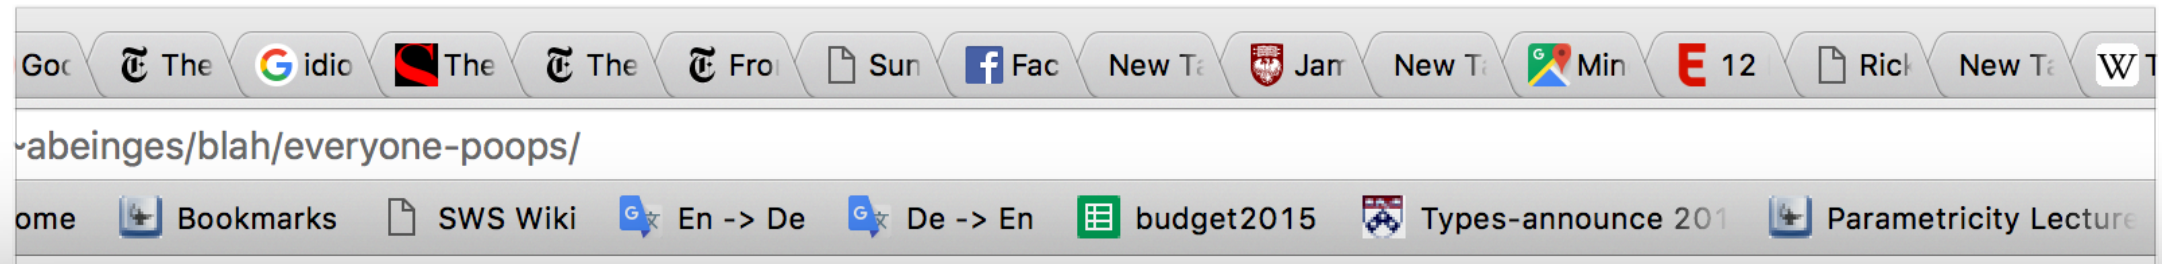
The issue in question states:

You can use a reference cycle to leak a JoinGuard and then the scoped thread can access freed memory

This is a very serious claim, since all the relevant APIs are marked as safe, and a use-after-free is something that should be *impossible* for safe code to perform.

The main focus is on the `thread::scoped` API which spawns a thread that can safely access the contents of another thread's stack frame *in a statically guaranteed way*. The basic idea is that `thread::scoped` returns a JoinGuard type. JoinGuard's destructor blocks on the thread joining, and isn't allowed to outlive any of the things that were passed into `thread::scoped`. This enables really nice things like:

Is Rust Safe?



Rust is at the **bleeding edge** of language design for safe systems programming

- We need **formal foundations** in order to build confidence in its safety guarantees!

memory

This is a very serious claim, since all the relevant APIs are marked as safe, and a use-after-free is something that should be *impossible* for safe code to perform.

The main focus is on the `thread::scoped` API which spawns a thread that can safely access the contents of another thread's stack frame *in a statically guaranteed way*. The basic idea is that `thread::scoped` returns a JoinGuard type. JoinGuard's destructor blocks on the thread joining, and isn't allowed to outlive any of the things that were passed into `thread::scoped`. This enables really nice things like:

RUSTBELT

Goal: **Develop 1st logical foundations for Rust**

- Use these foundations to verify the safety of the Rust core type system and std libraries
- Give Rust developers the tools they need to safely evolve the language

What is “Safety”?


What is “Safety”?

- Standard “**syntactic safety**” approach of Wright and Felleisen (1994) **will not work for Rust!**
 - Requires whole program to be well-typed!


What is “Safety”?

- Standard “**syntactic safety**” approach of Wright and Felleisen (1994) **will not work for Rust!**
 - Requires whole program to be well-typed!
- Need to generalize to **semantic safety**
 - A library is semantically safe if no well-typed application using it can have undefined behavior

Semantic Safety



Library
interface


semantic
model




Safety
contract

\models
logical
satisfaction




Library
implementation

Semantic Safety




Library
interface


semantic
model




Safety
contract


logical
satisfaction


Well
typed

Library
implementation

Semantic Safety


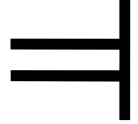


Library
interface


semantic
model




Safety
contract



logical
satisfaction




Library
implementation

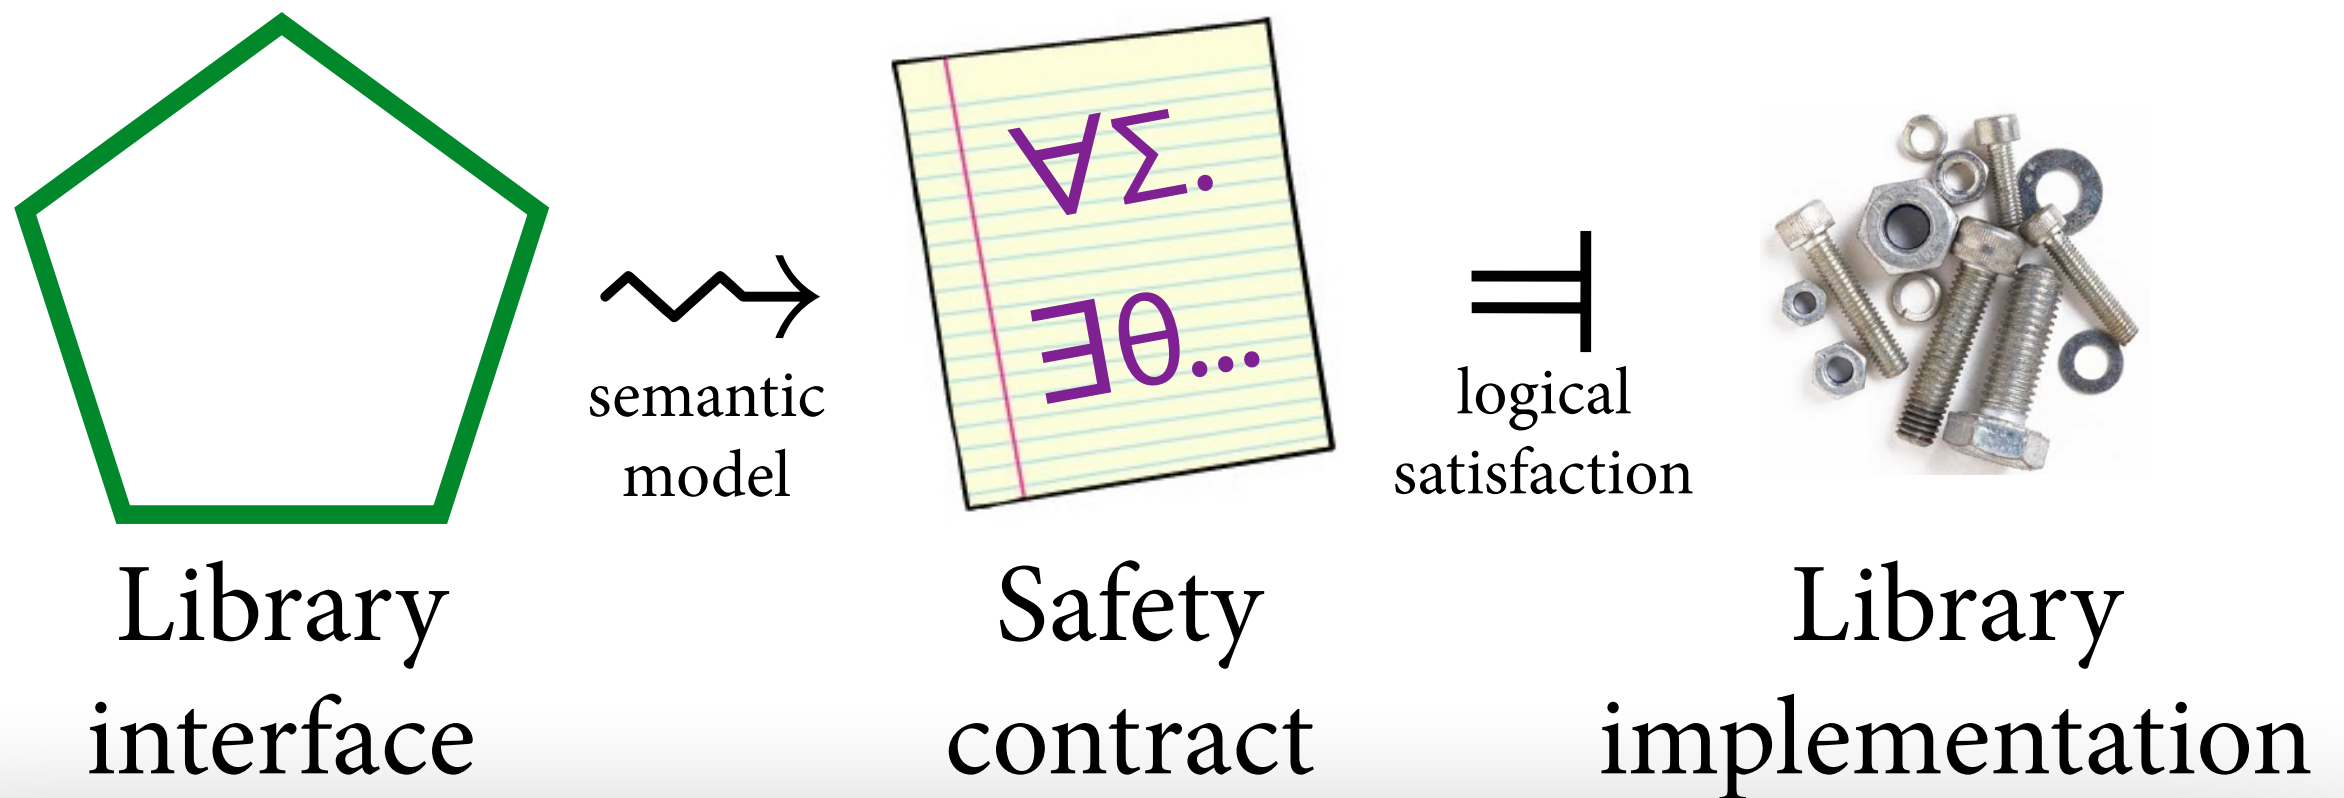
Semantic Safety



Semantic Safety



Semantic Safety



Challenge:

Verify semantic safety for



Heart of the Problem



Which logic to use?

Separation
Logic

to the
Rescue!



Separation Logic

to the
Rescue!



Extension of Hoare logic (O'Hearn, Reynolds..., ~2000)

- For reasoning about pointer-manipulating programs



Major influence on many verification & analysis tools

- e.g. Infer, VeriFast, Chalice, Bedrock, jStar, ...



Separation logic = Ownership logic

- Perfect fit for modeling Rust's ownership types!

Problem 1: Which One?



Problem 1: Which One?





Concurrent Separation Logic
[O'Hearn/Brookes, 2007]



Won the 2016 Gödel Prize!

Problem 1: Which One?



Problem 1: Which One?



The Next 700 Separation Logics (Invited Paper)

Matthew Parkinson

Microsoft Research Cambridge

Abstract. In recent years, separation logic has brought great advances in the world of verification. However, there is a disturbing trend for each new library or concurrency primitive to require a new separation logic. I will argue that we shouldn't be inventing new separation logics, but should find the right logic to reason about interference, and have a powerful abstraction mechanism to enable the library's implementation details to be correctly abstracted. Adding new concurrency libraries should simply be a matter of verification, not of new logics or metatheory.

Landin's seminal paper, The Next 700 Programming Languages [33], opens with:

Most programming languages are partly a way of expressing things in terms of other things and partly a basic set of given things.



Problem 1: Which One?



The Next 700 Separation Logics (Invited Paper)


Matthew Parkinson

Microsoft Research Cambridge

Abstract. In recent years, separation logic has brought great advances in the world of verification. However, there is a disturbing trend for each new library or concurrency primitive to require a new separation logic. I will argue that we shouldn't be inventing new separation logics, but should find the right logic to reason about interference, and have a powerful abstraction mechanism to enable the library's implementation details to be correctly abstracted. Adding new concurrency libraries should simply be a matter of verification, not of new logics or metatheory.

Landin's seminal paper, The Next 700 Programming Languages [33], opens with:

Most programming languages are partly a way of expressing things in terms of other things and partly a basic set of given things.



Problem 2: Memory Model



All these logics assume:

- **sequential consistency** for memory accesses

*This is totally **unrealistic** for high-performance concurrency!*

- e.g. Rust's Arc library uses **C++'s weak memory ops**

Towards a Logic for Rust

- **Iris** [POPL'15, ICFP'16, POPL'17, ESOP'17]:
Simplifying & unifying modern separation logics
+ support for machine-checked proof in Coq
- **GPS** [OOPSLA'14, PLDI'15, ECOOP'17]:
First modern sep. logic for C++ memory model

In these lectures...

- Day 1: Ownership types
- Day 2: Concurrent separation logic
- Day 3: Introduction to Iris framework
& how we are using it to verify safety of Rust

In these lectures



- Day 1: Ownership types **in Rust!**
- Day 2: Concurrent separation logic **in Coq!**
- Day 3: **Interactive demos!** newwork
& how we are using it to verify safety of Rust