



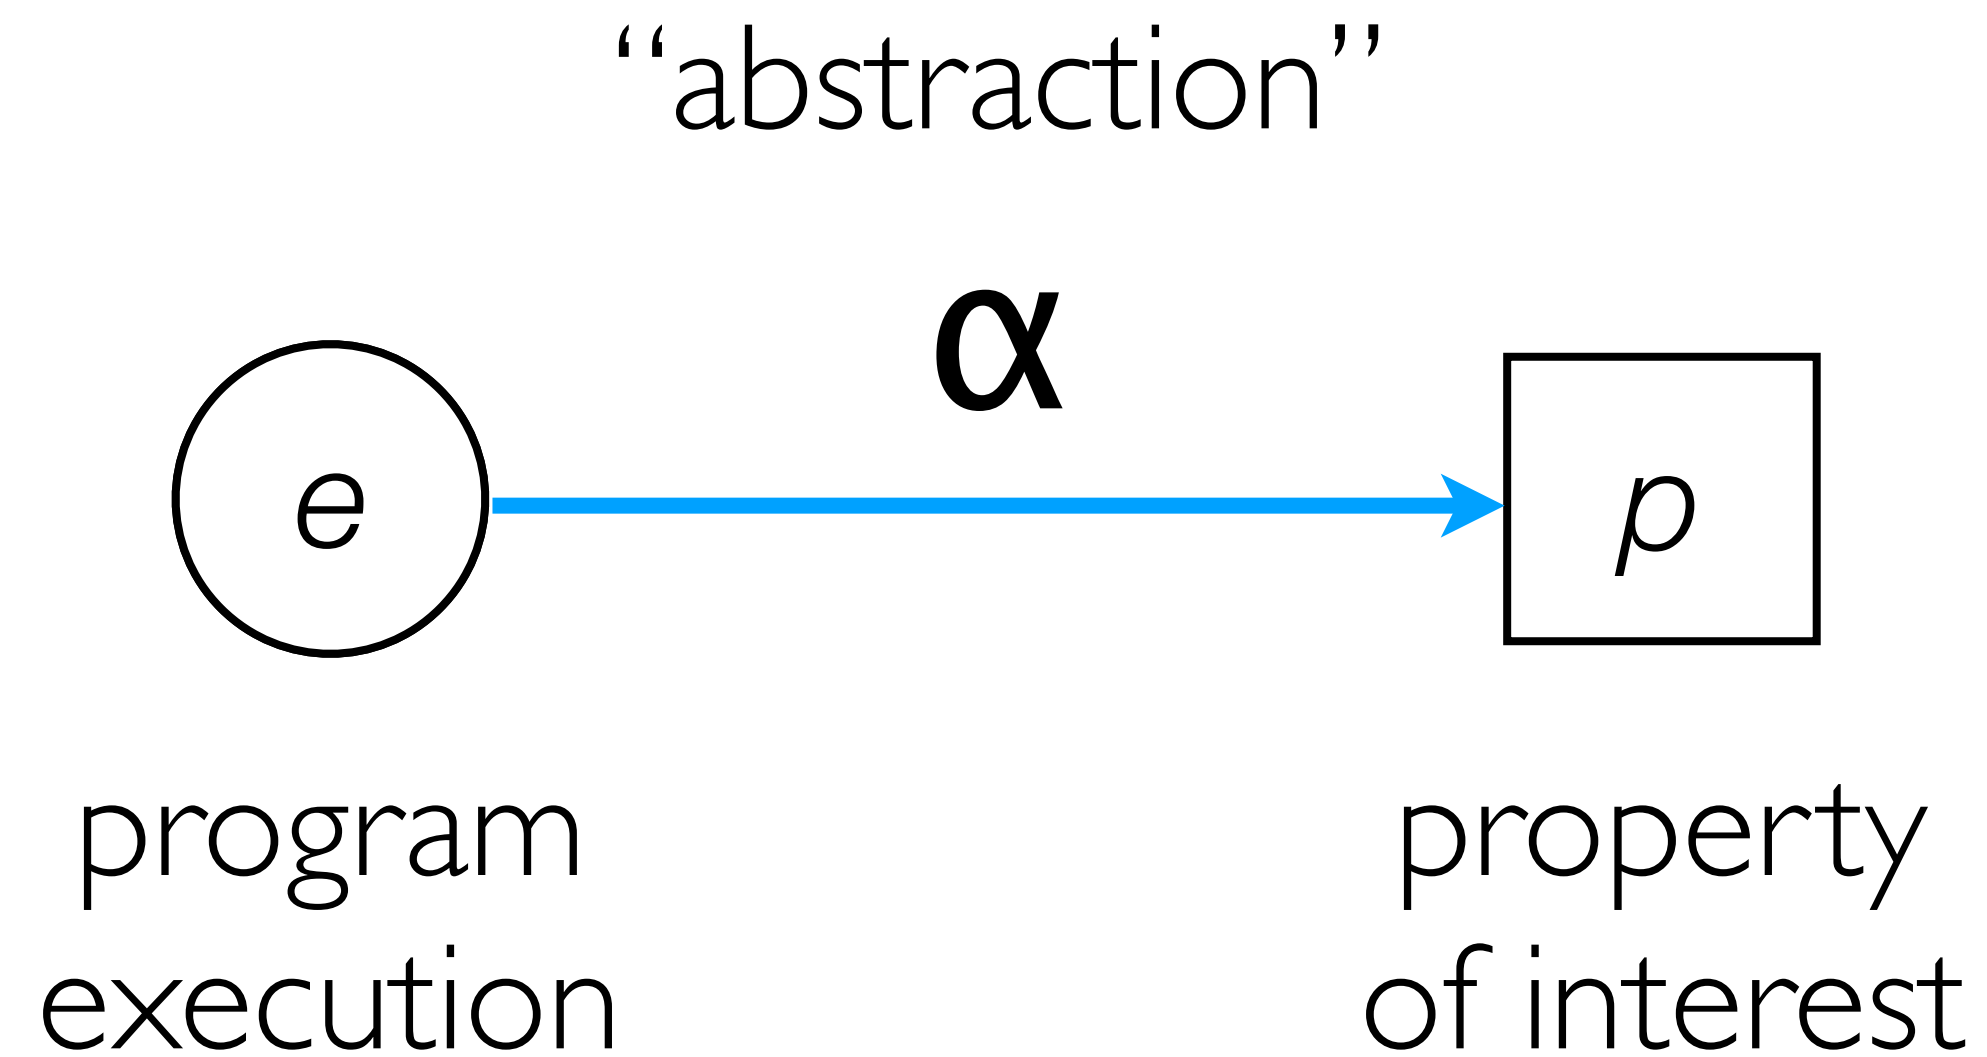
A True Positives Theorem for a Static Race Detector

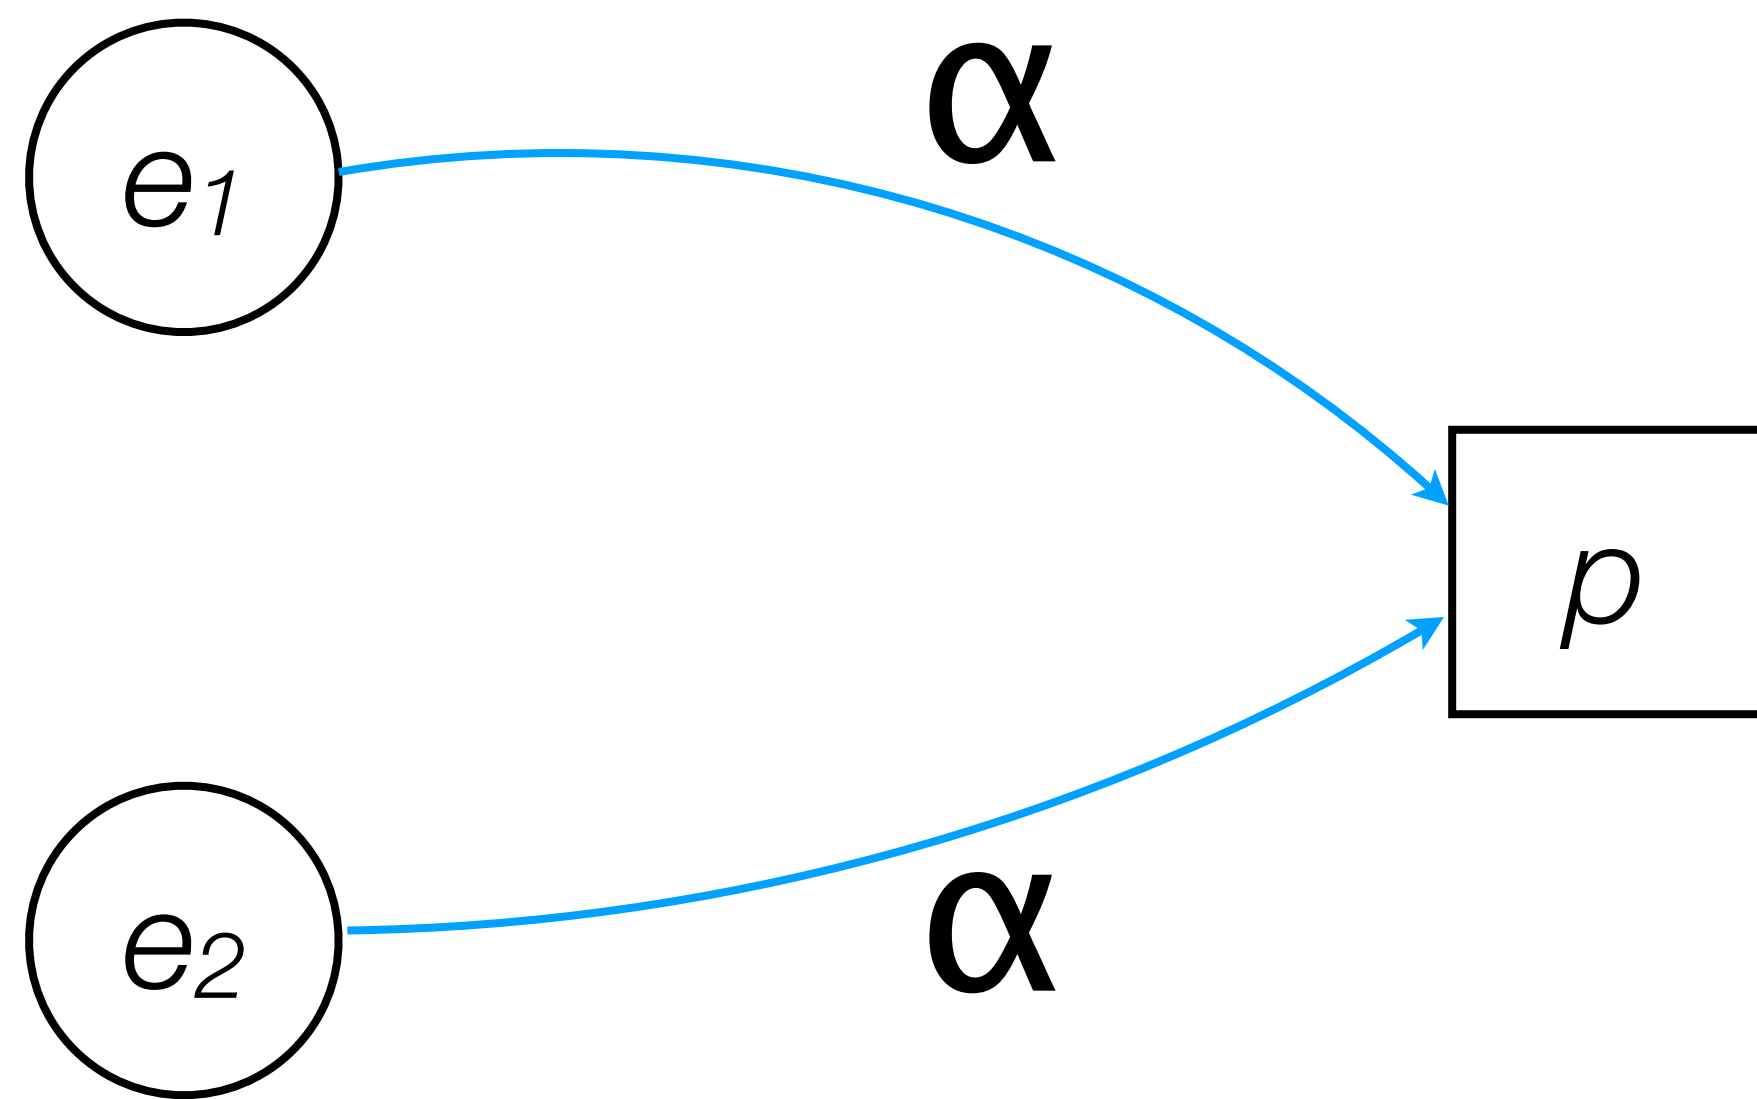
Presentation by Julia Belyakova and Artem Pelenitsyn
For [CS 7580](#) (instructor: Jan Vitek), 10/30/2019

A subset of slides is taken from [Ilya Sergey's web page](#)

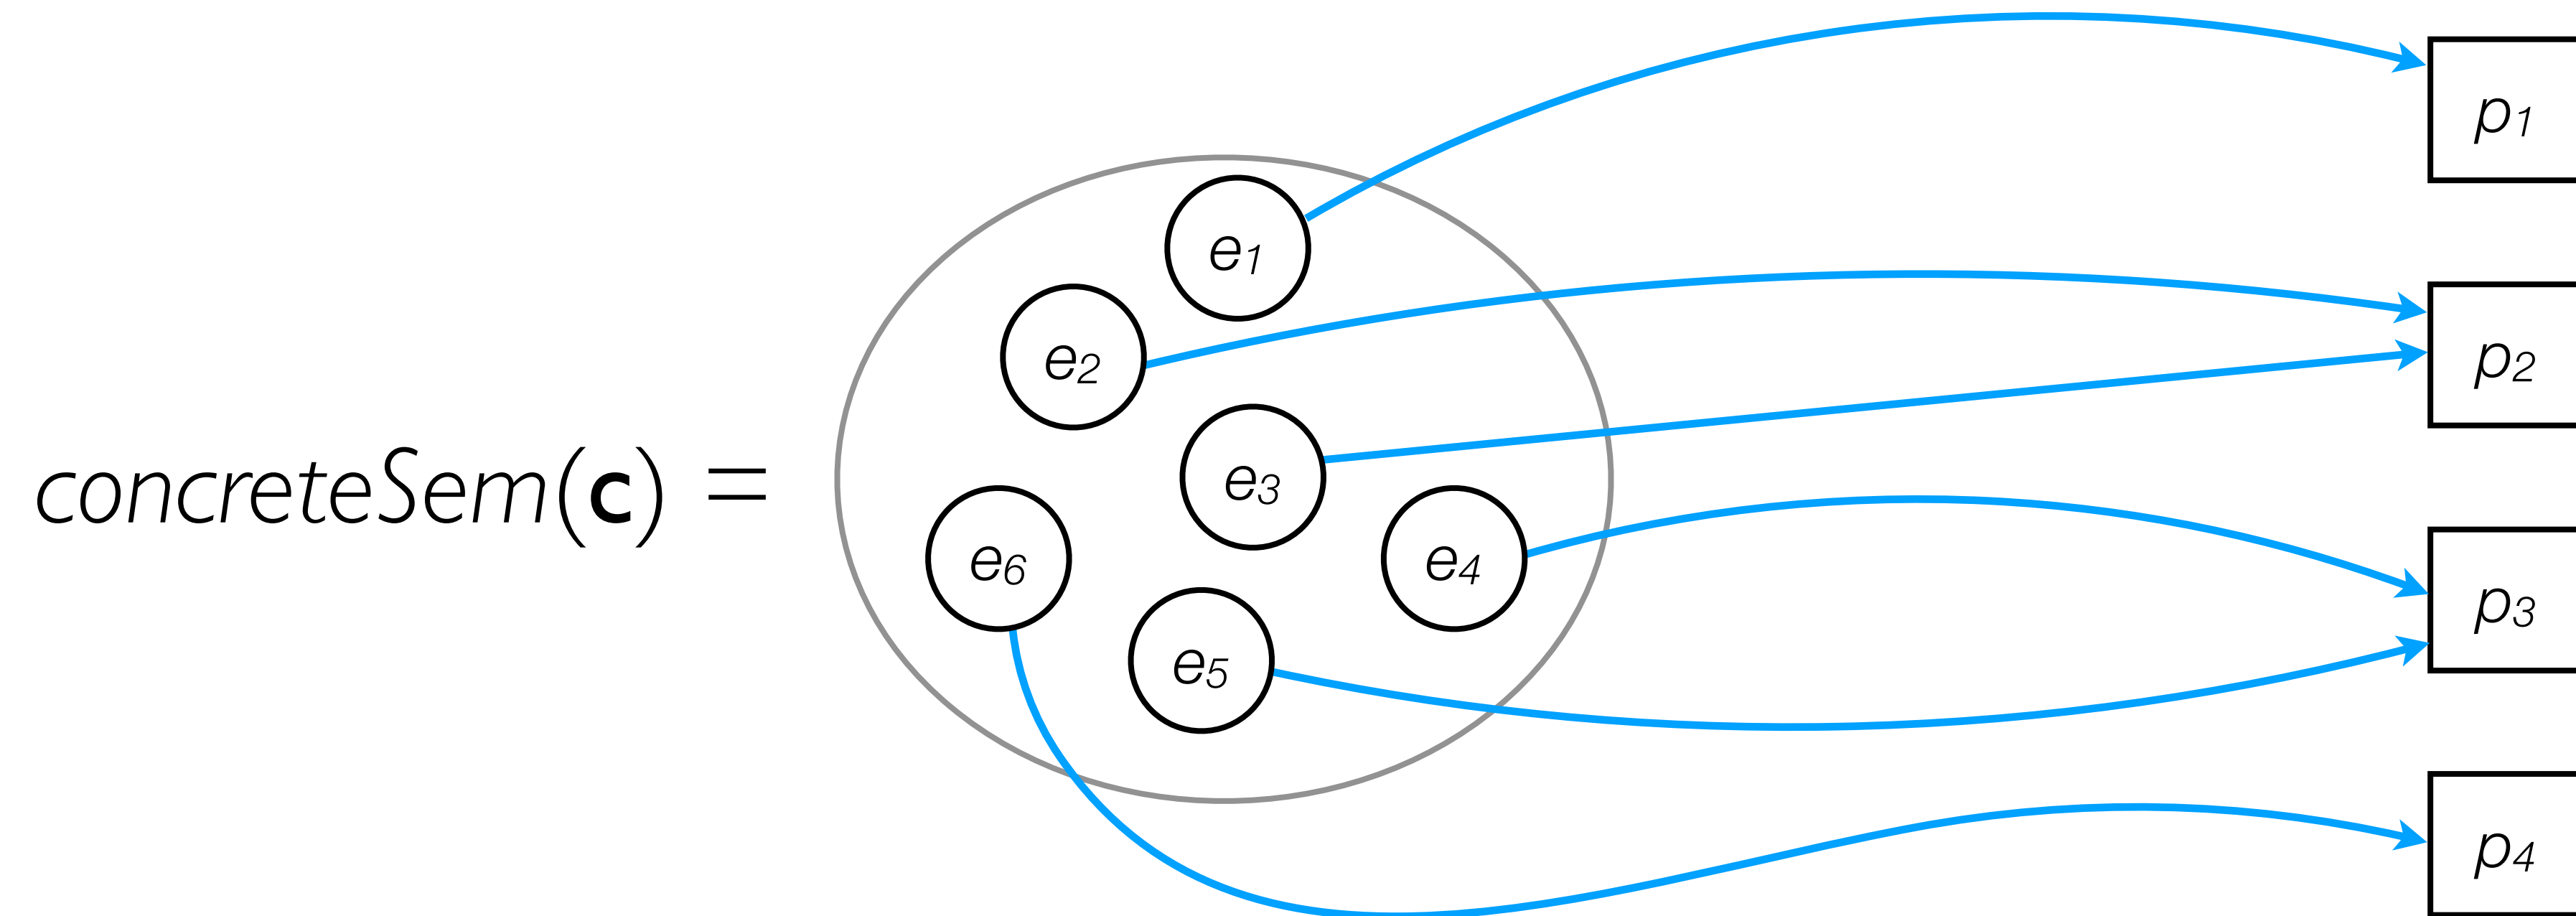
Static Analyses for Program Validation

The Essence of Static Analysis

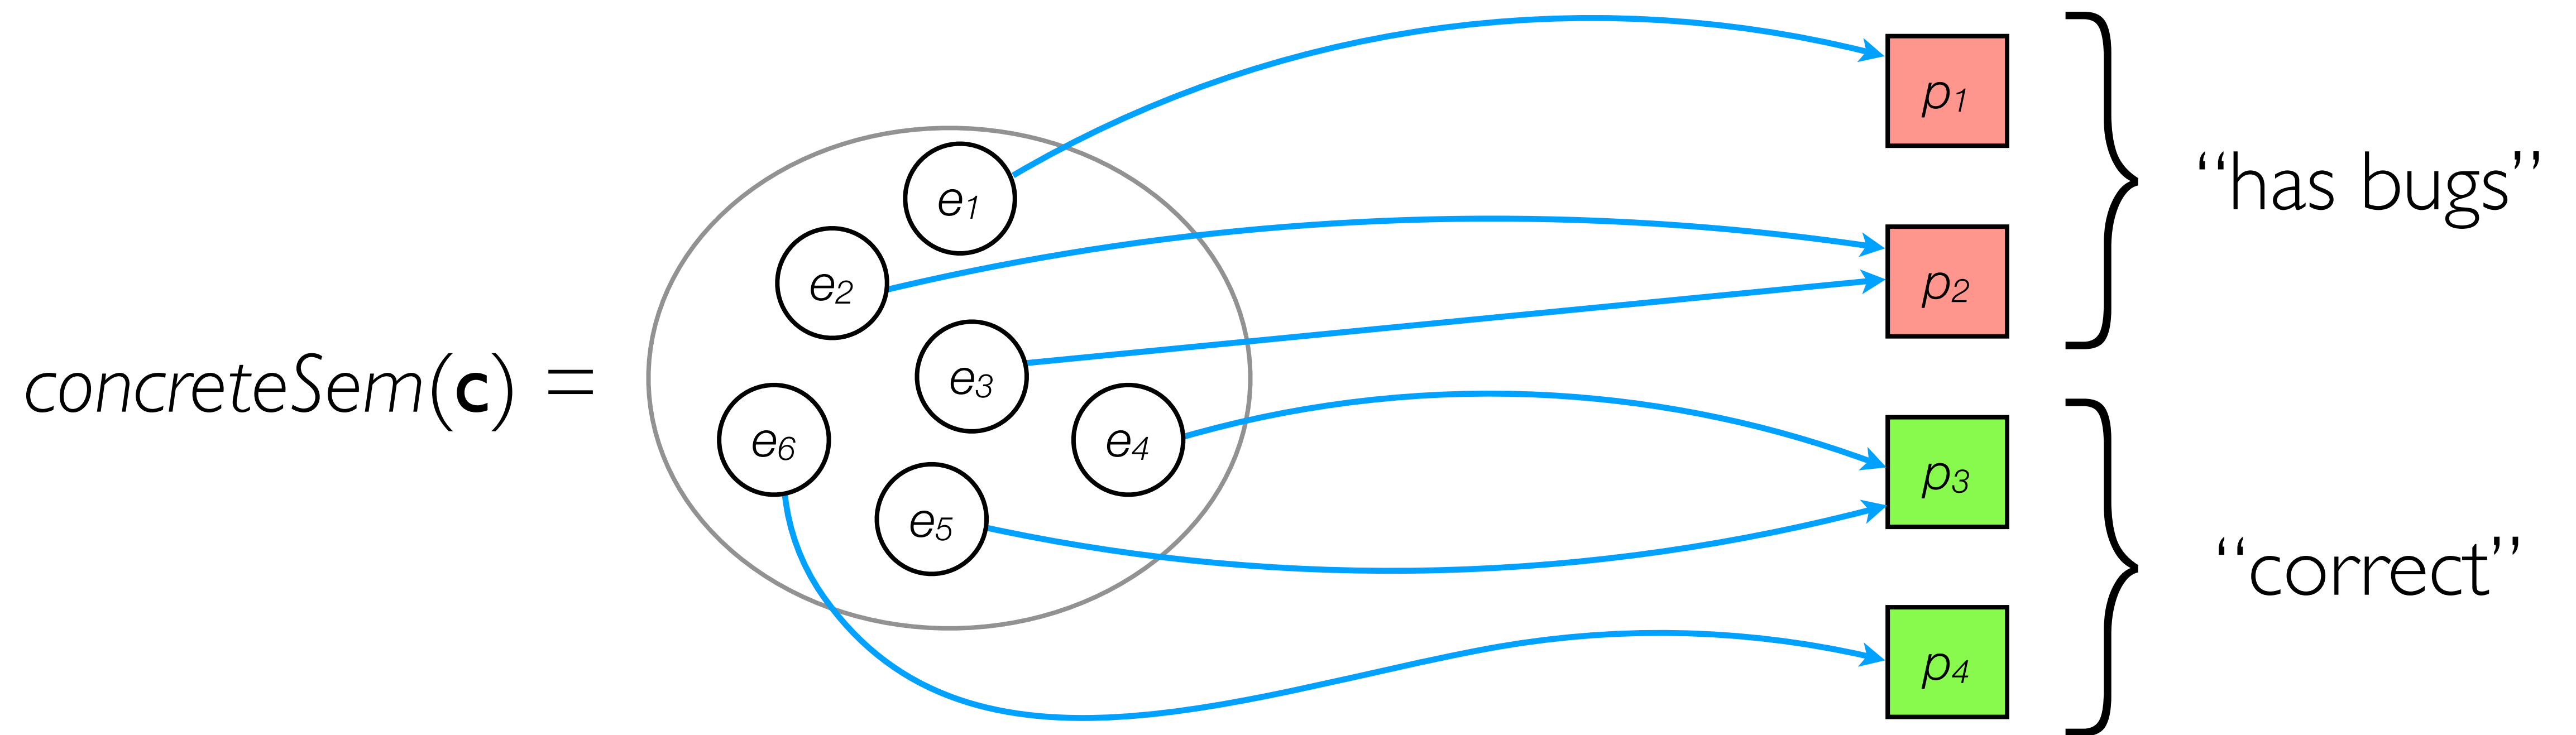




Static Analysis

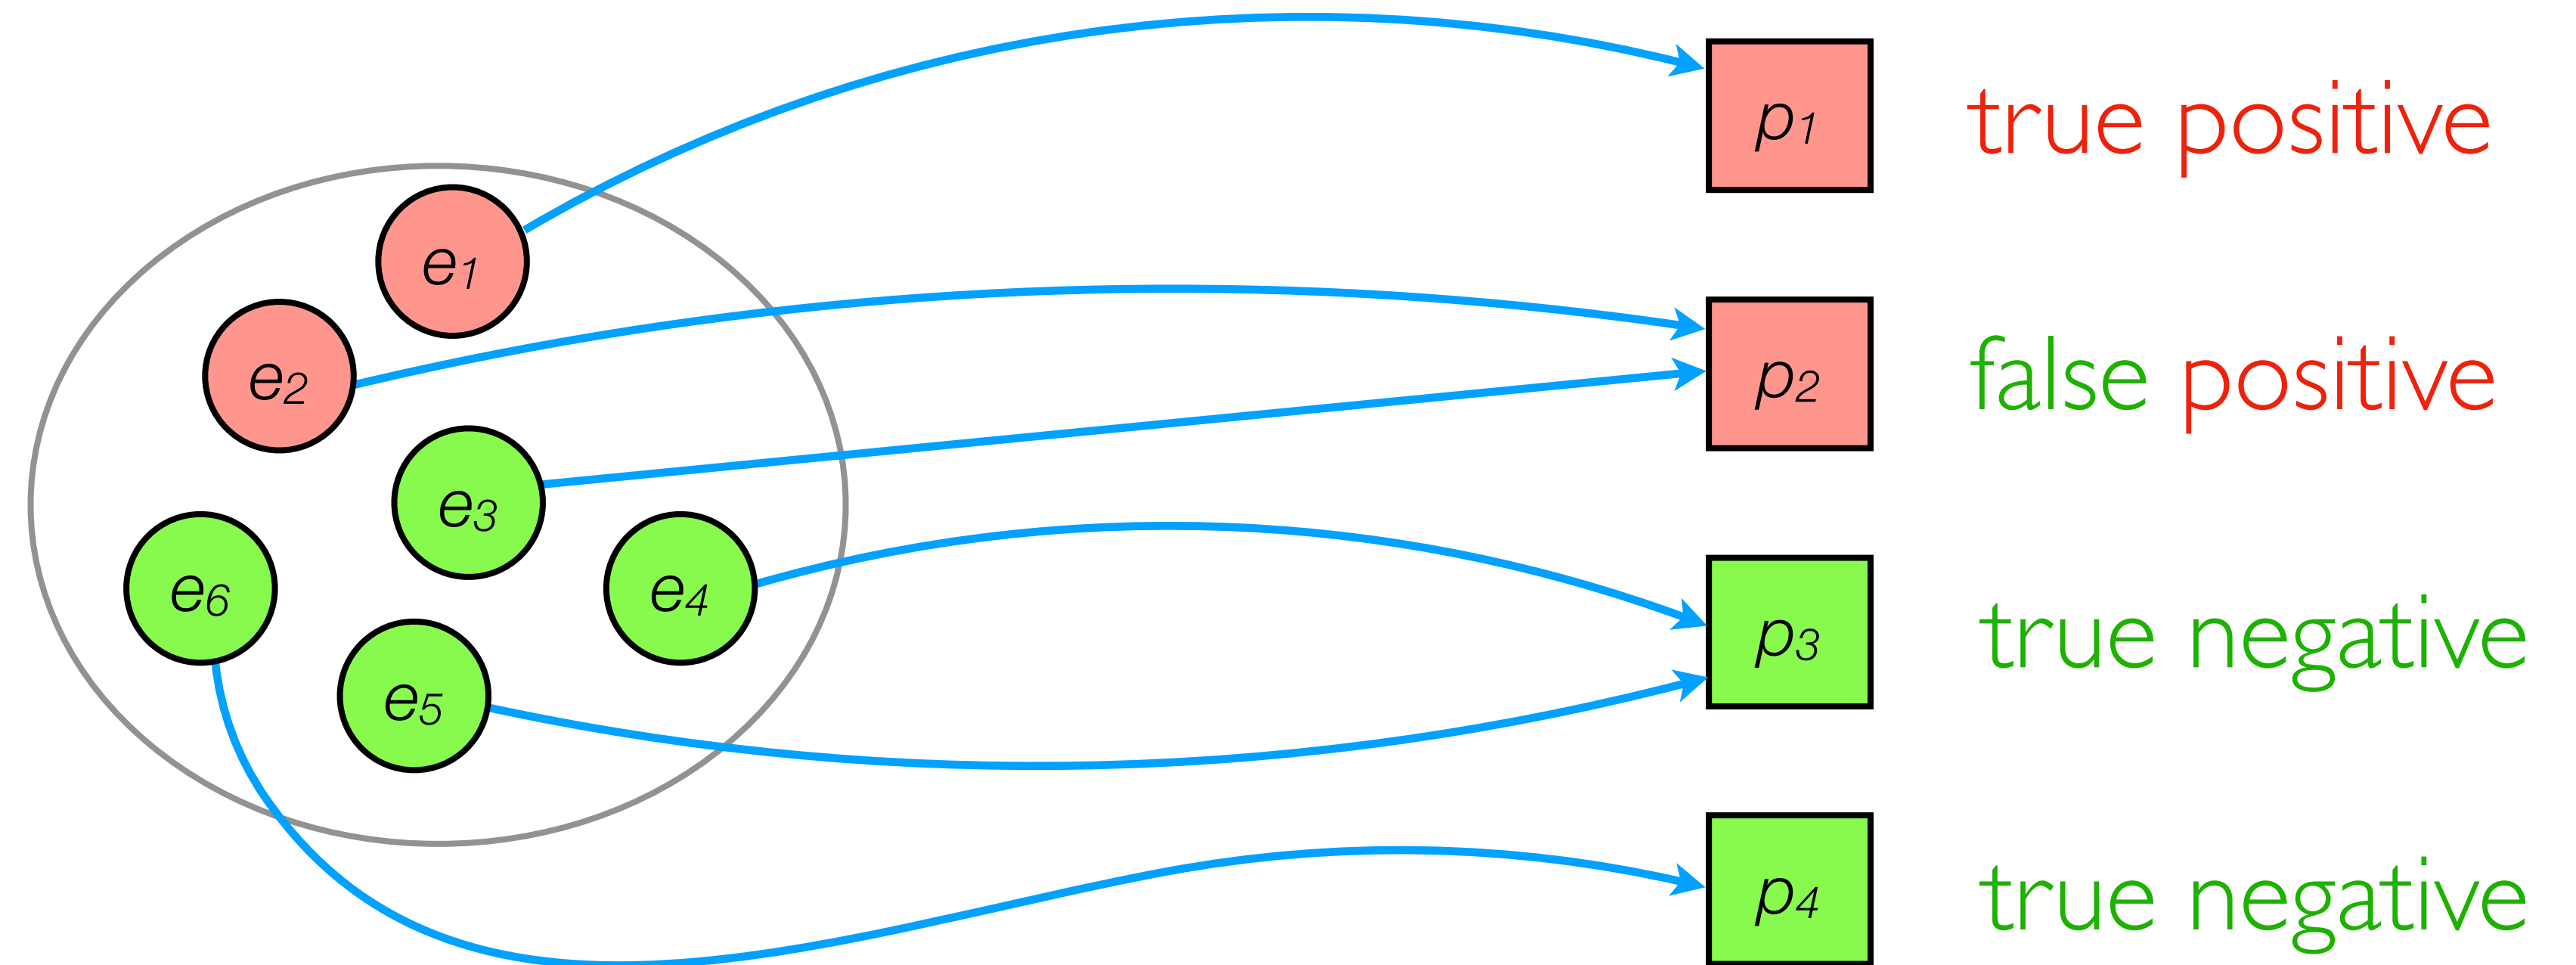


Static Analysis

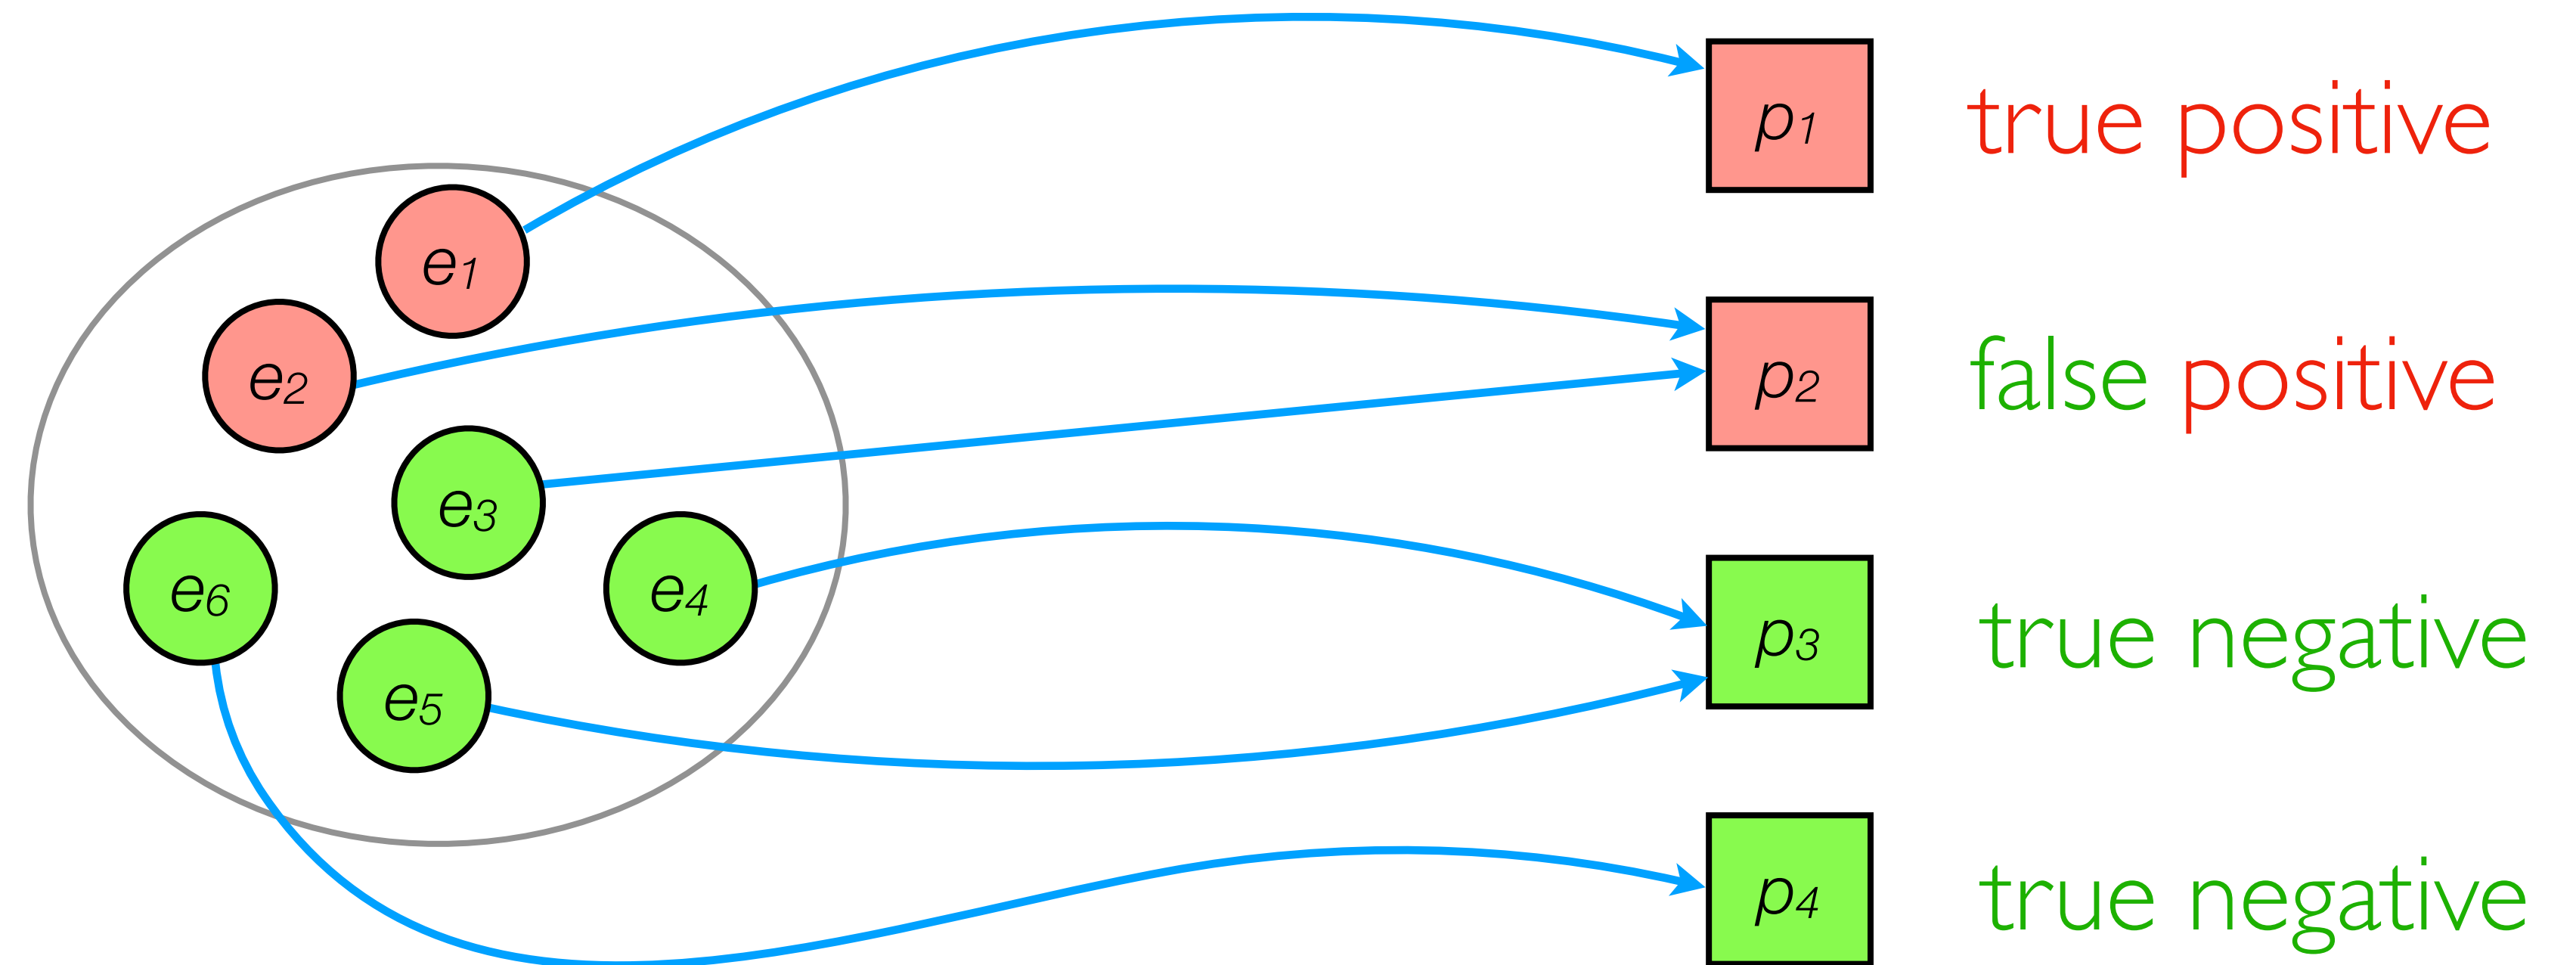


Verifier
or a
Bug Detector?

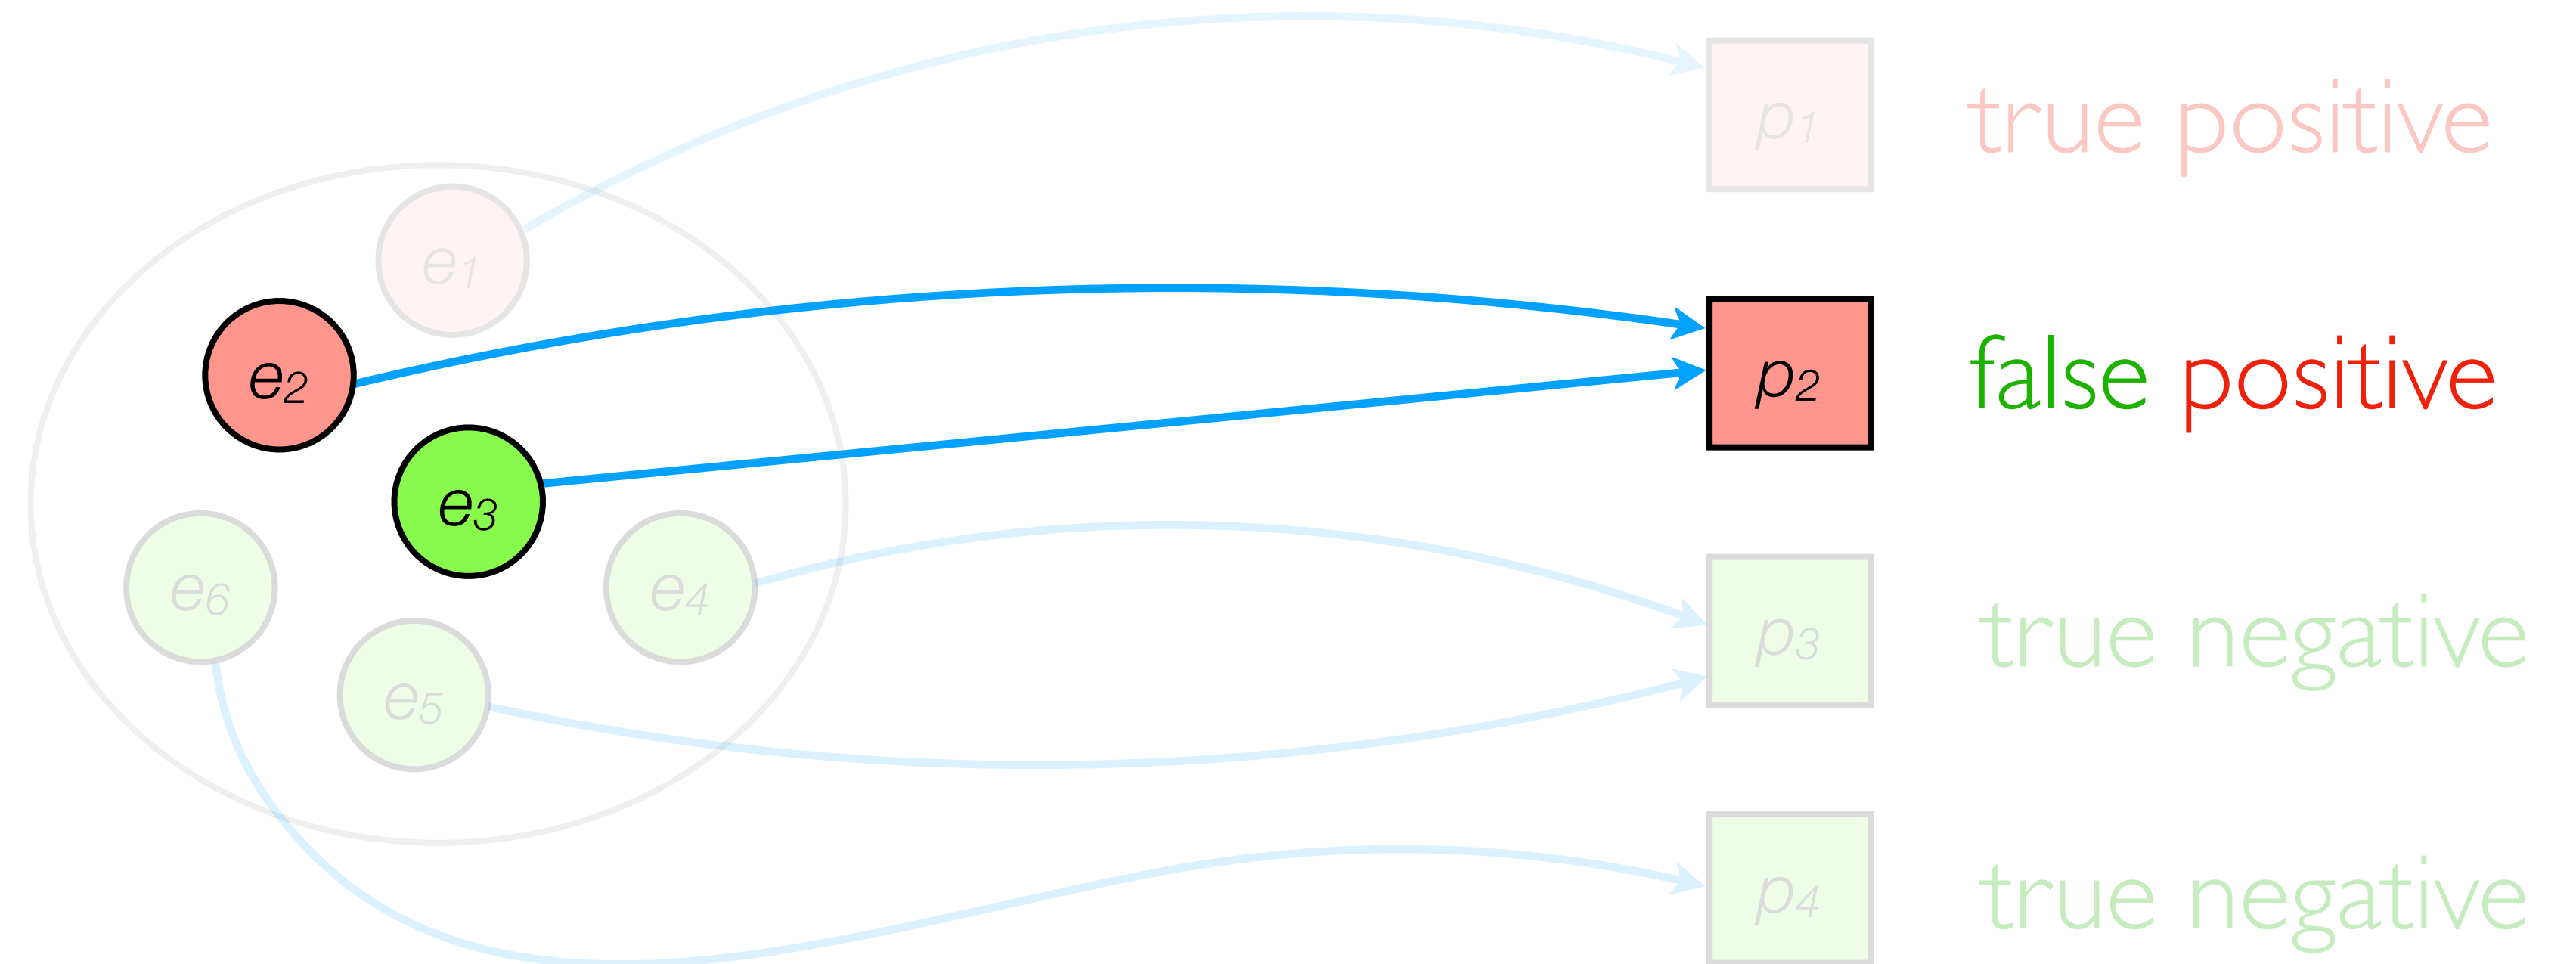
Program Verifier



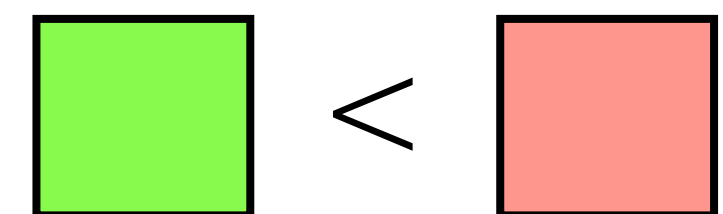
Sound Program Verifier



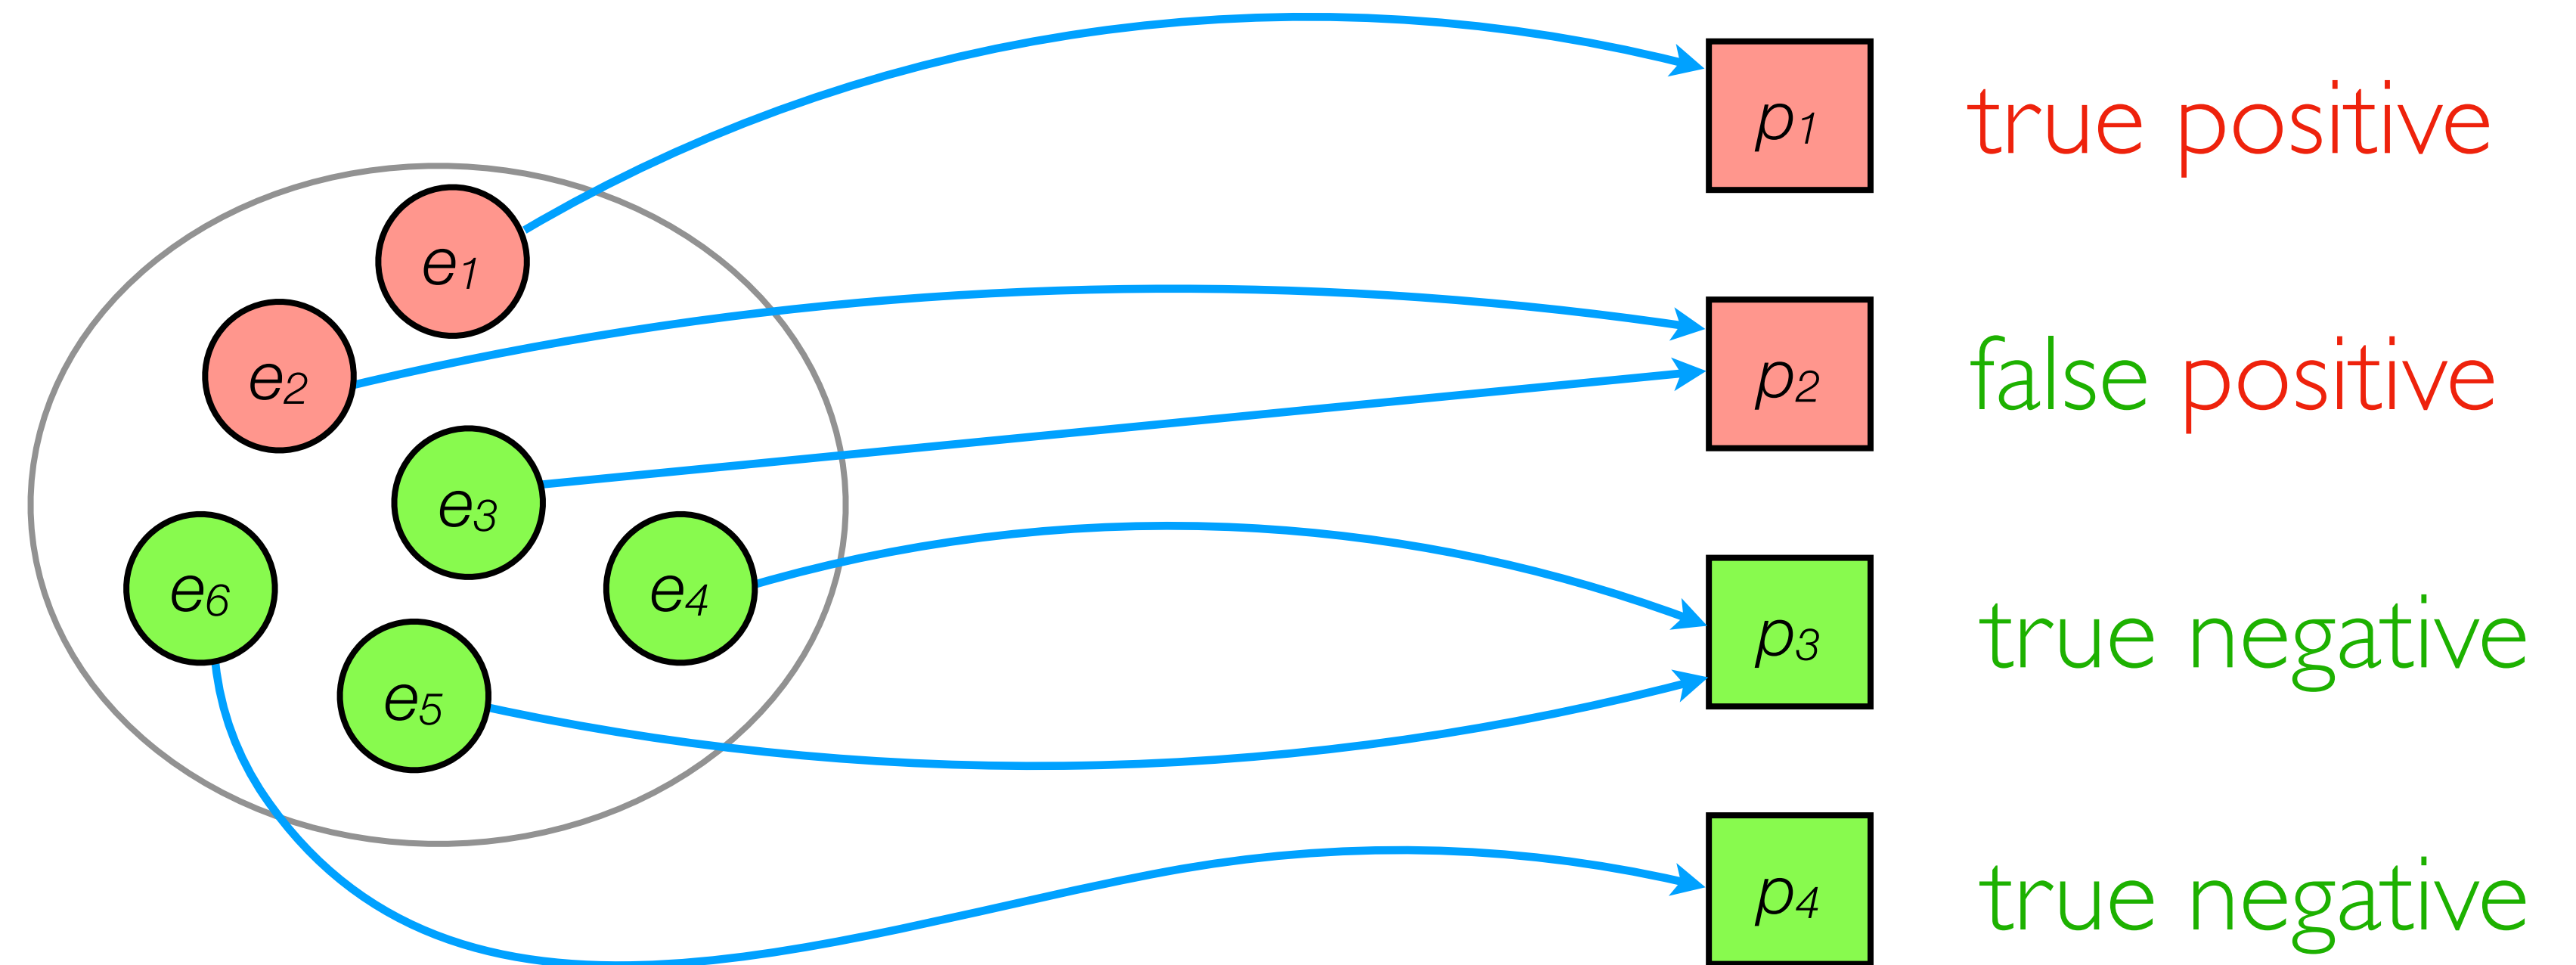
Sound Program Verifier



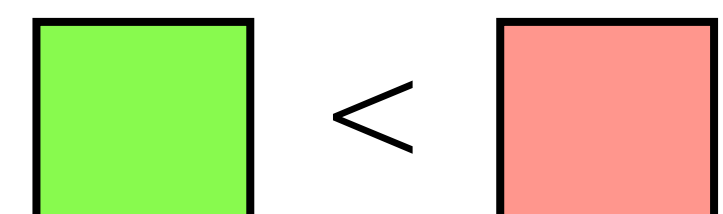
abstract over-approximation



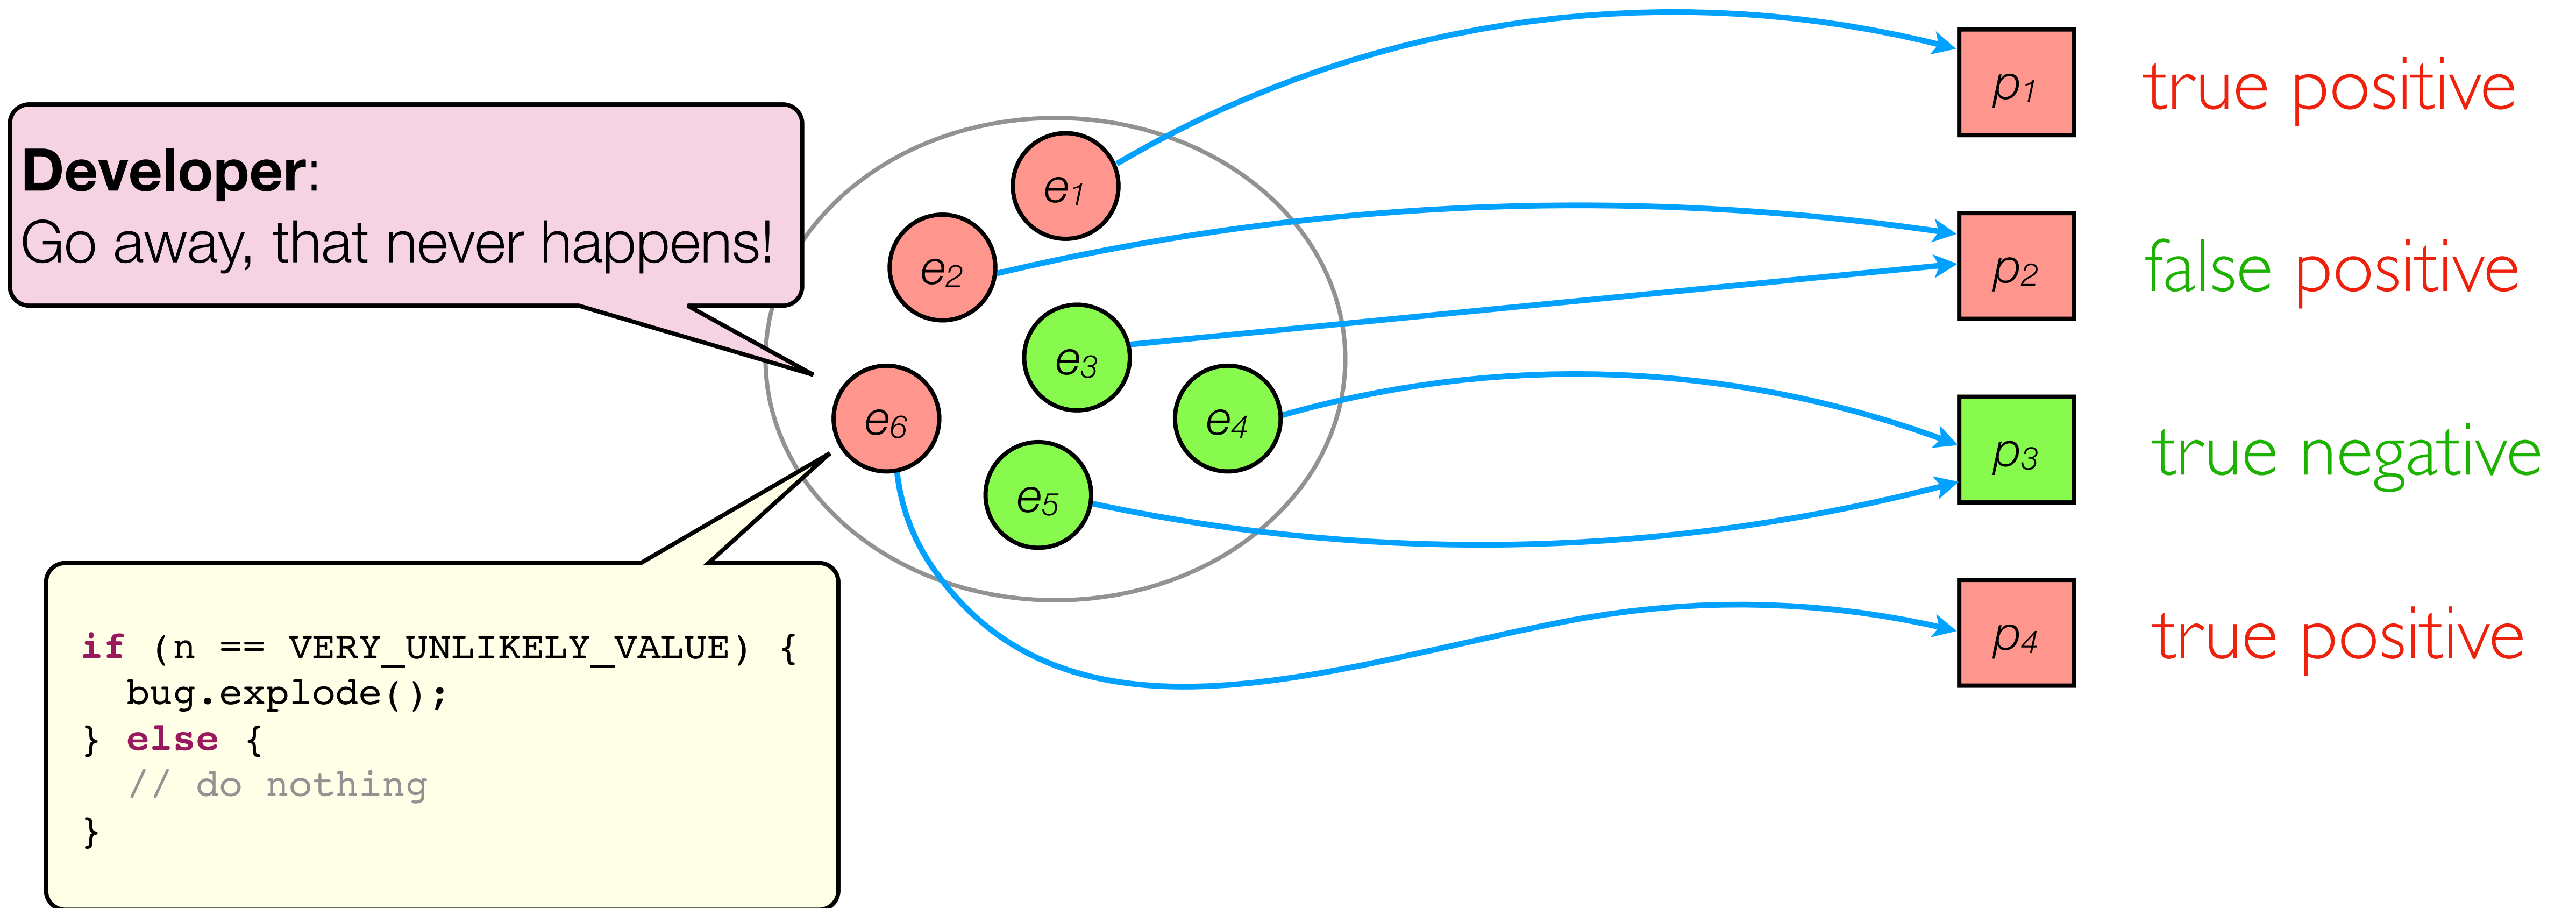
Sound Program Verifier



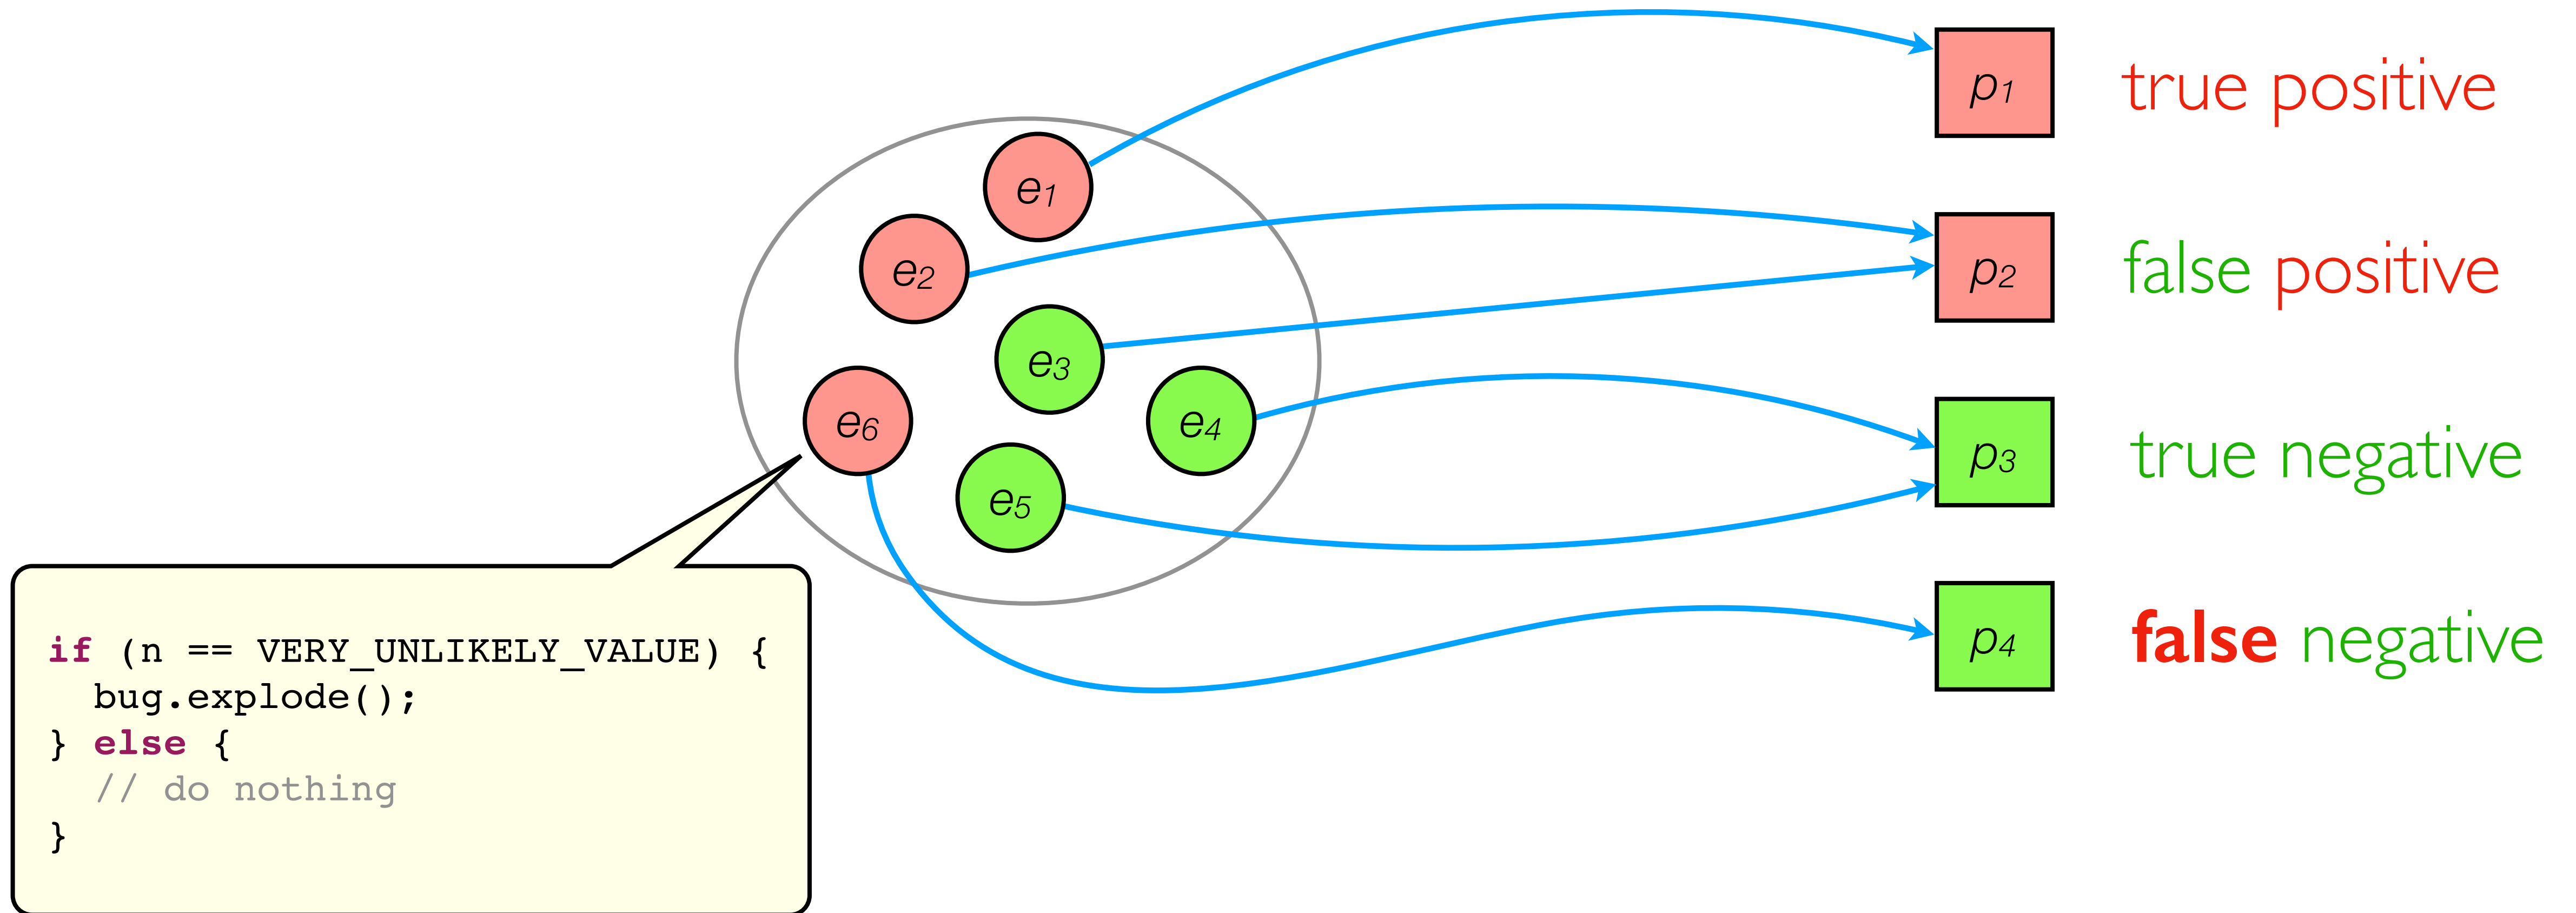
abstract over-approximation



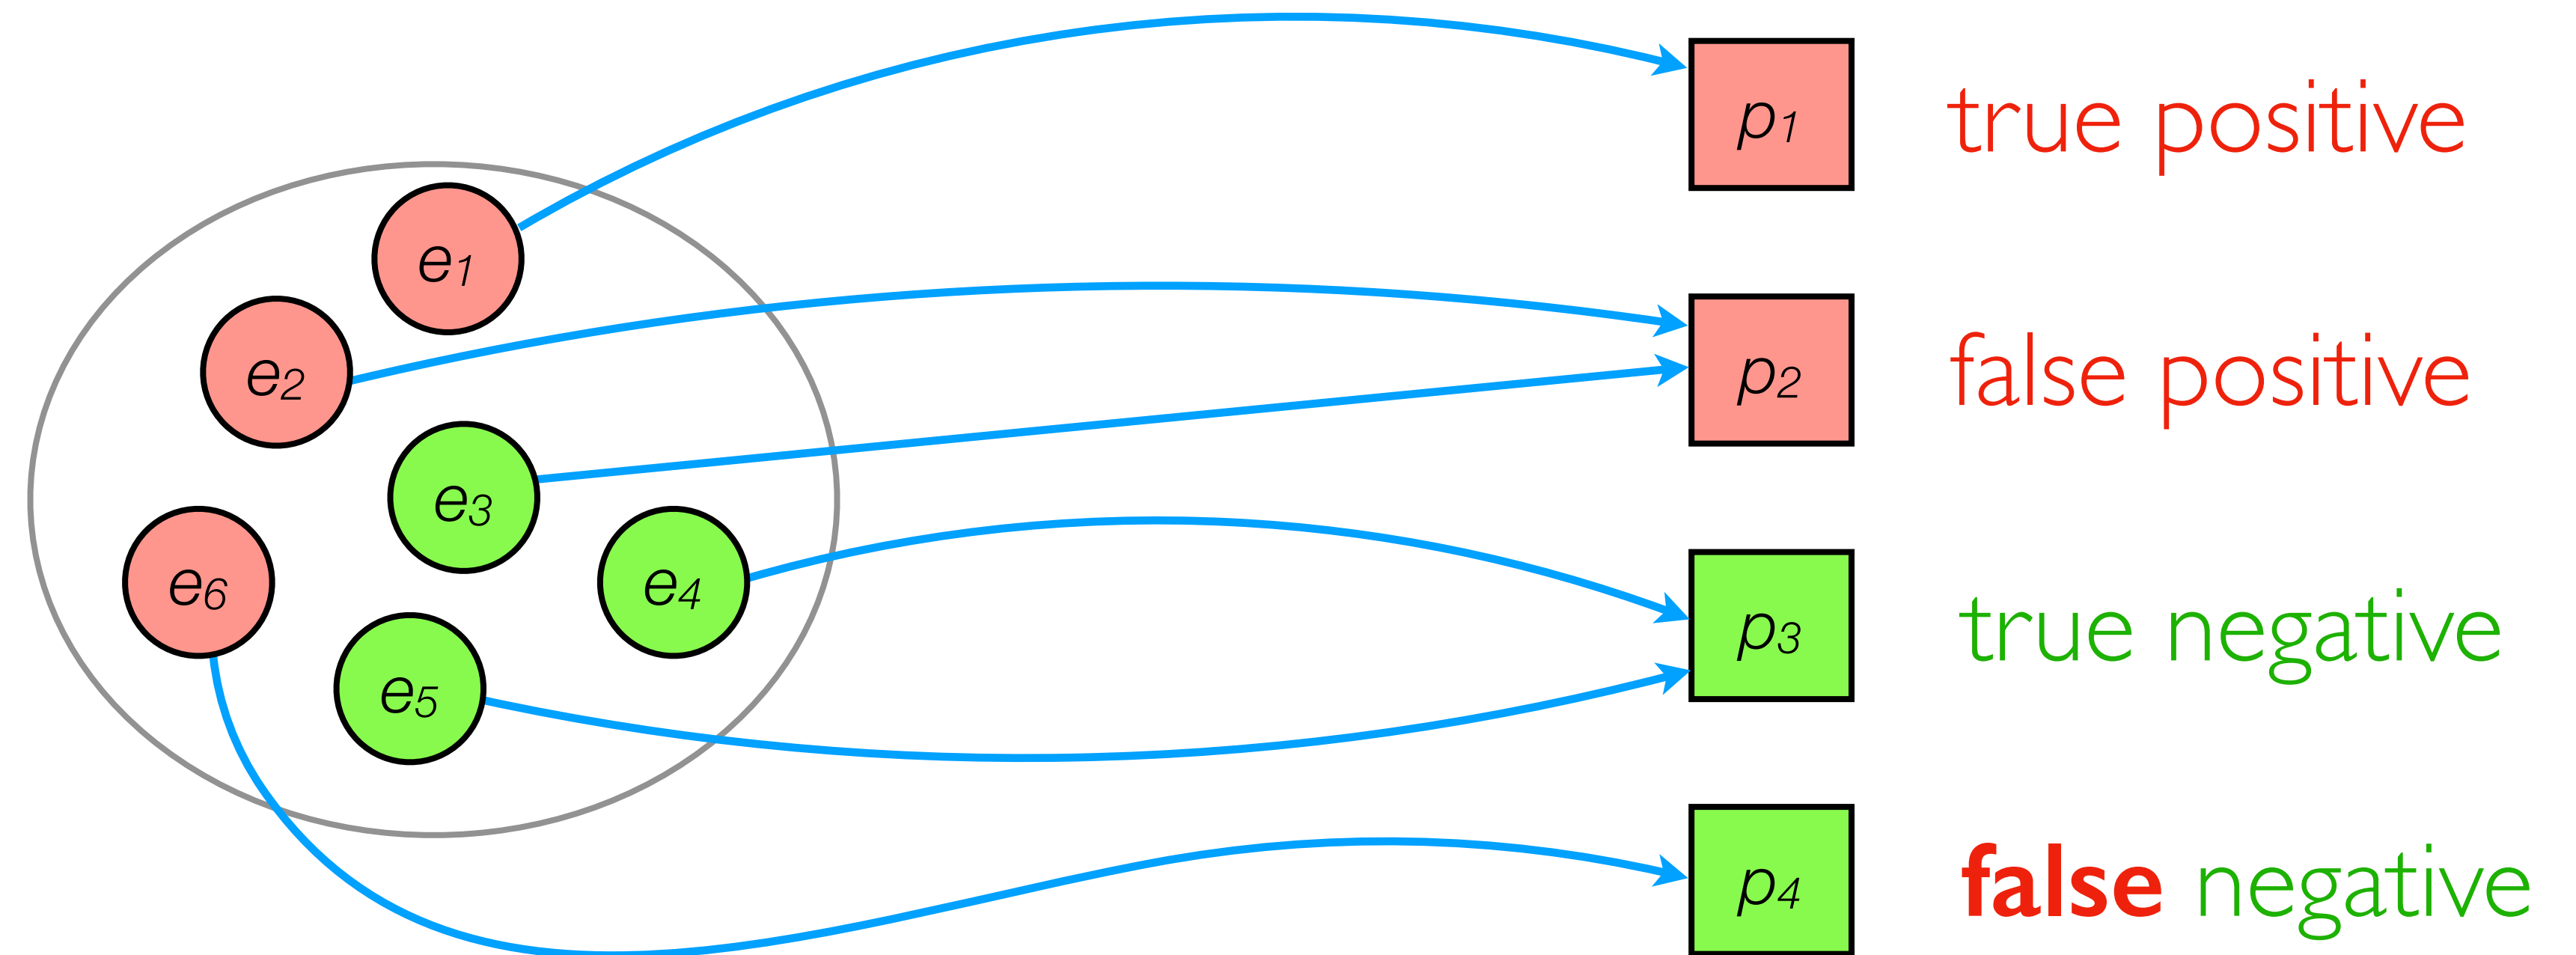
Sound Program Verifier



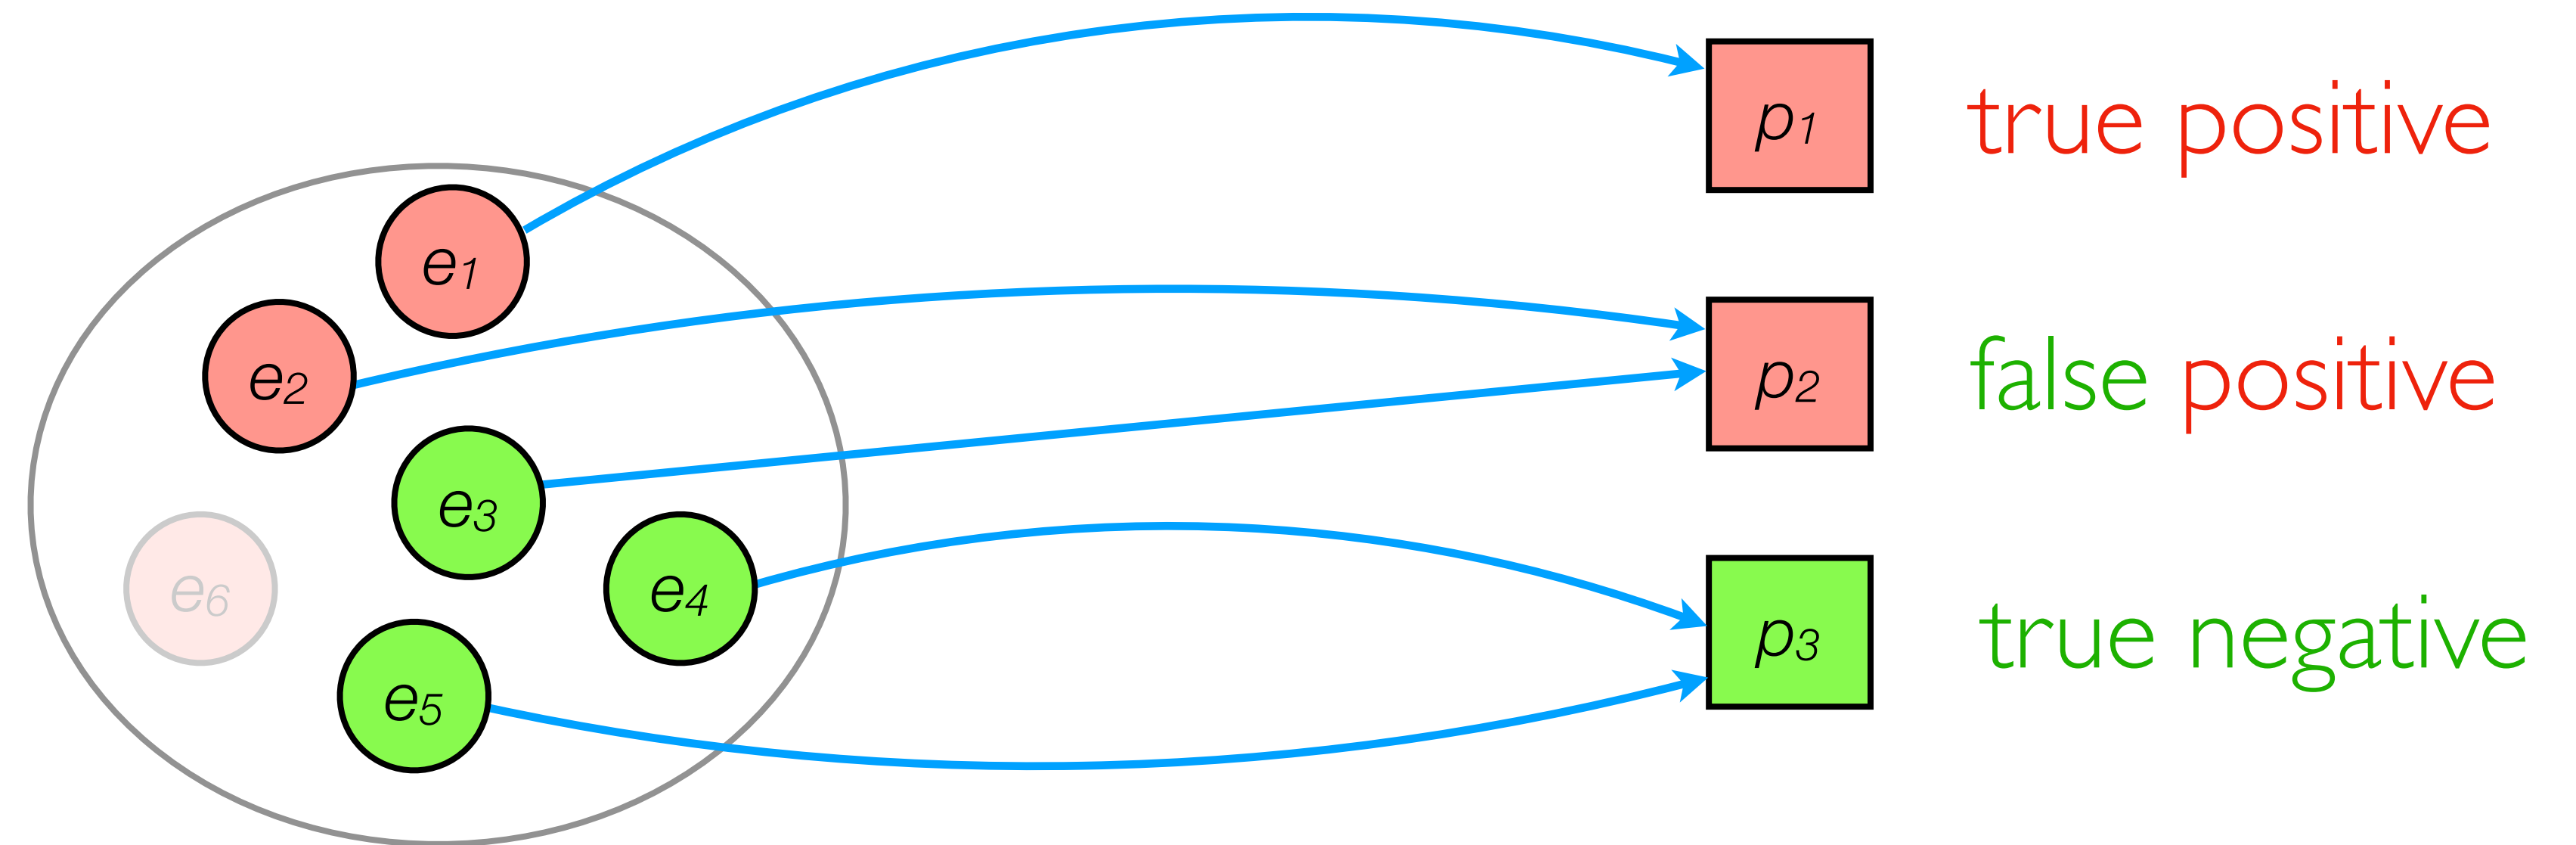
Unsound Program “Verifier”



“Sound” Program Verifier

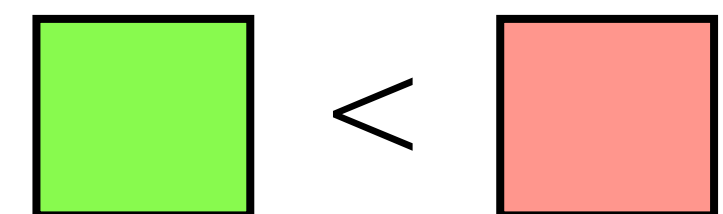


“Sound” Program Verifier



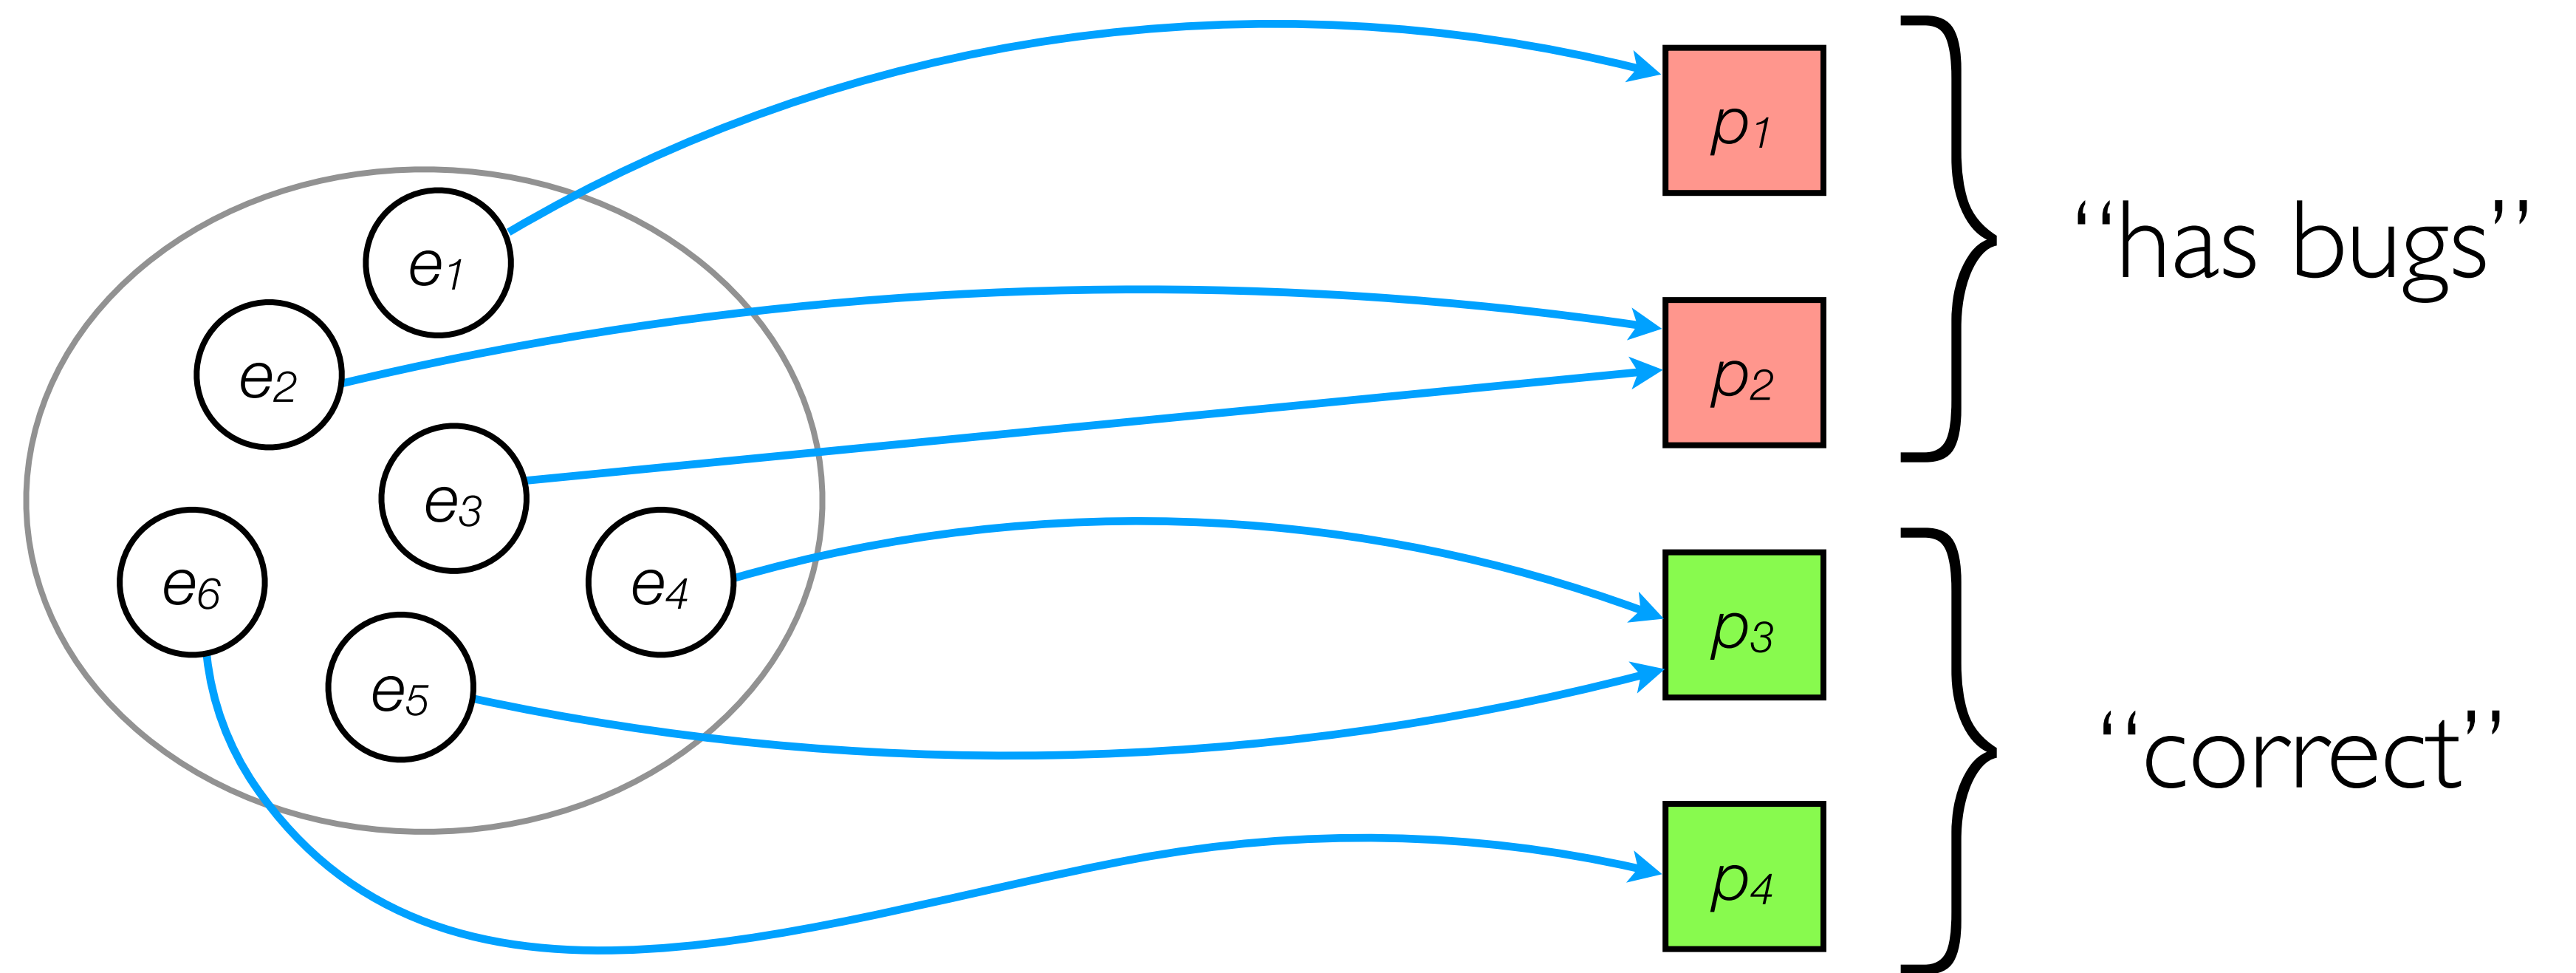
concrete under-approximation

abstract over-approximation

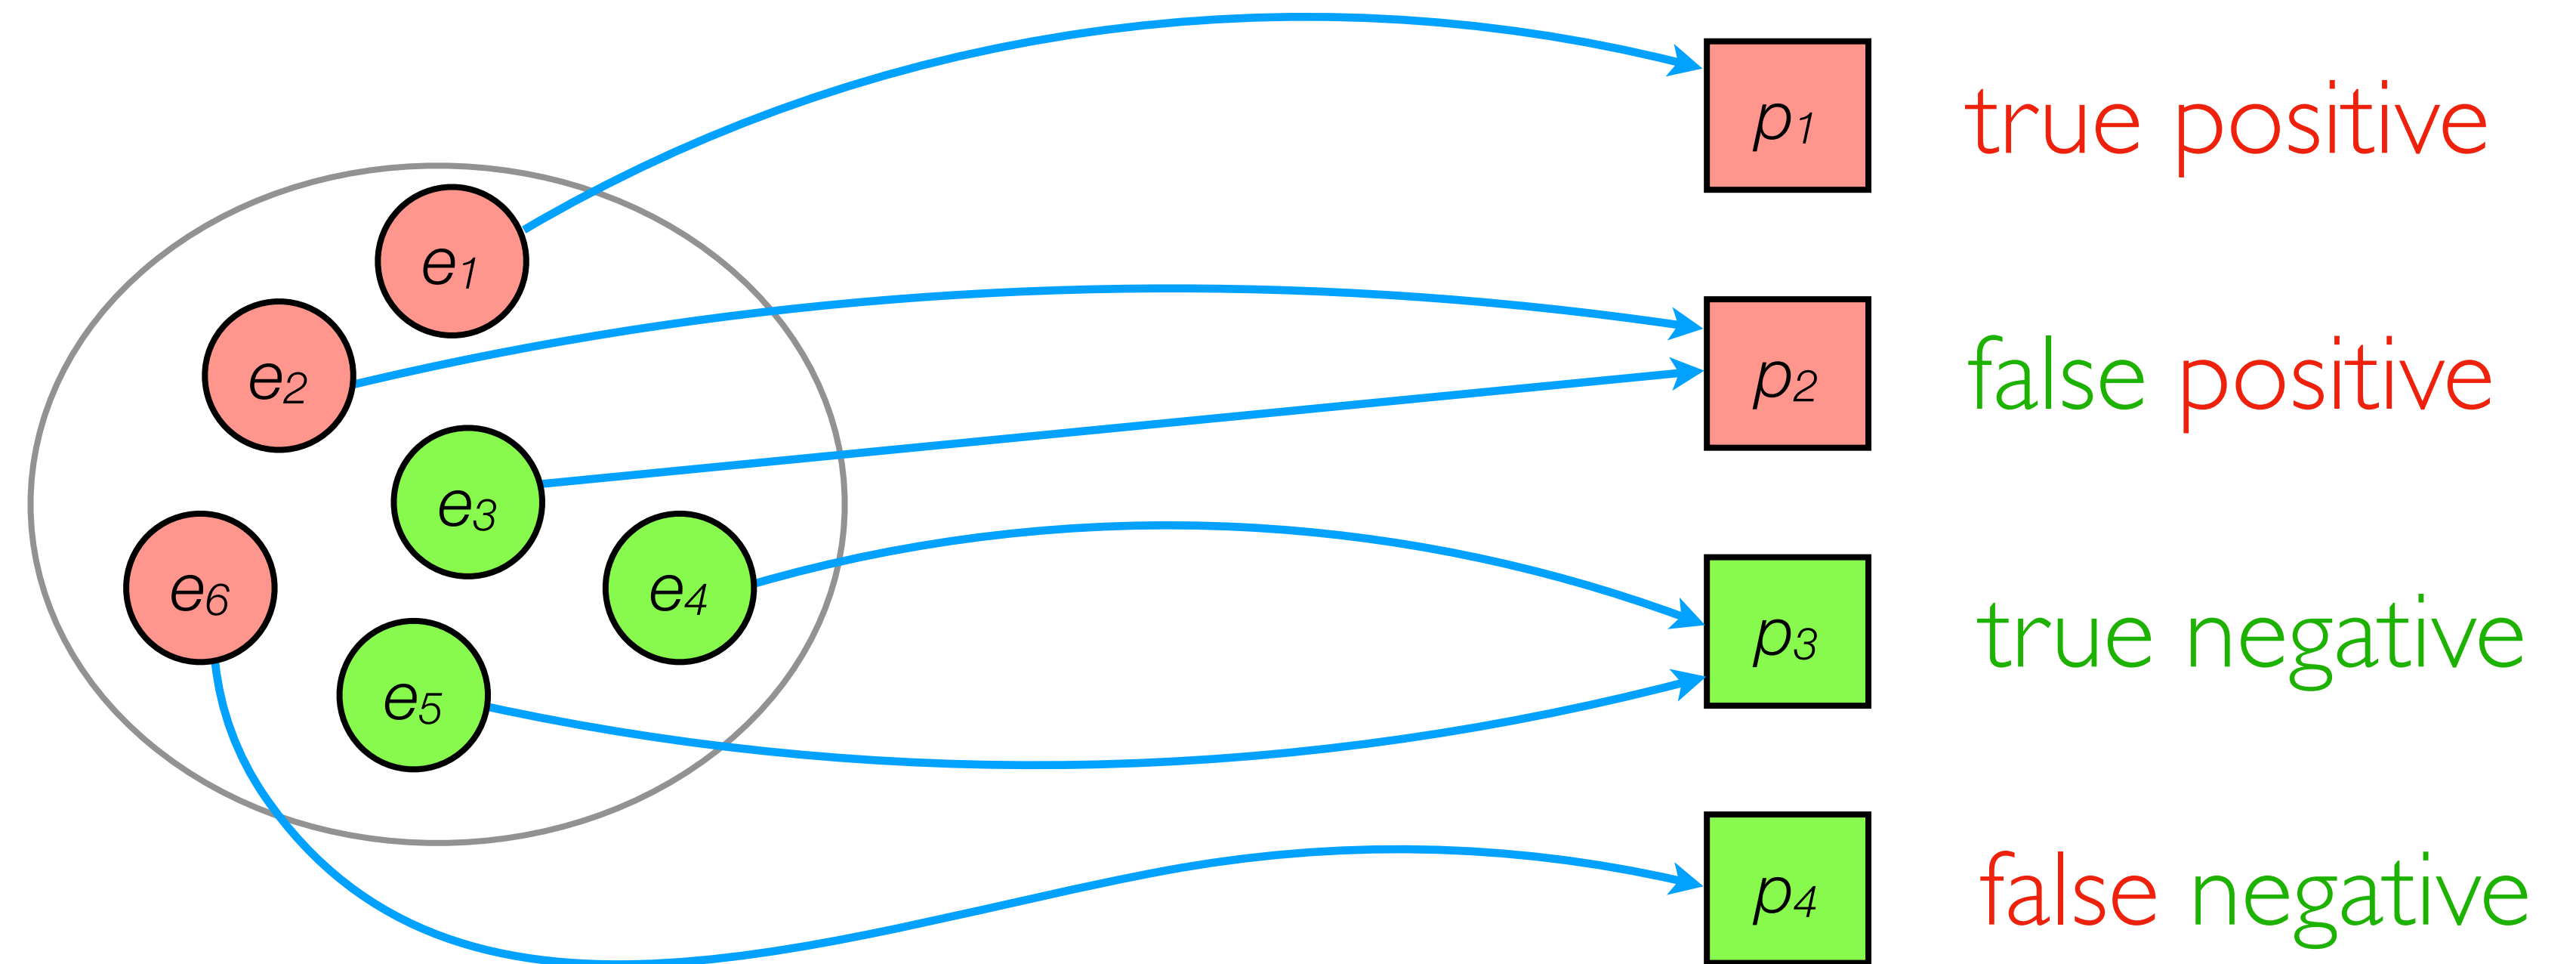


Sound Static Verifiers

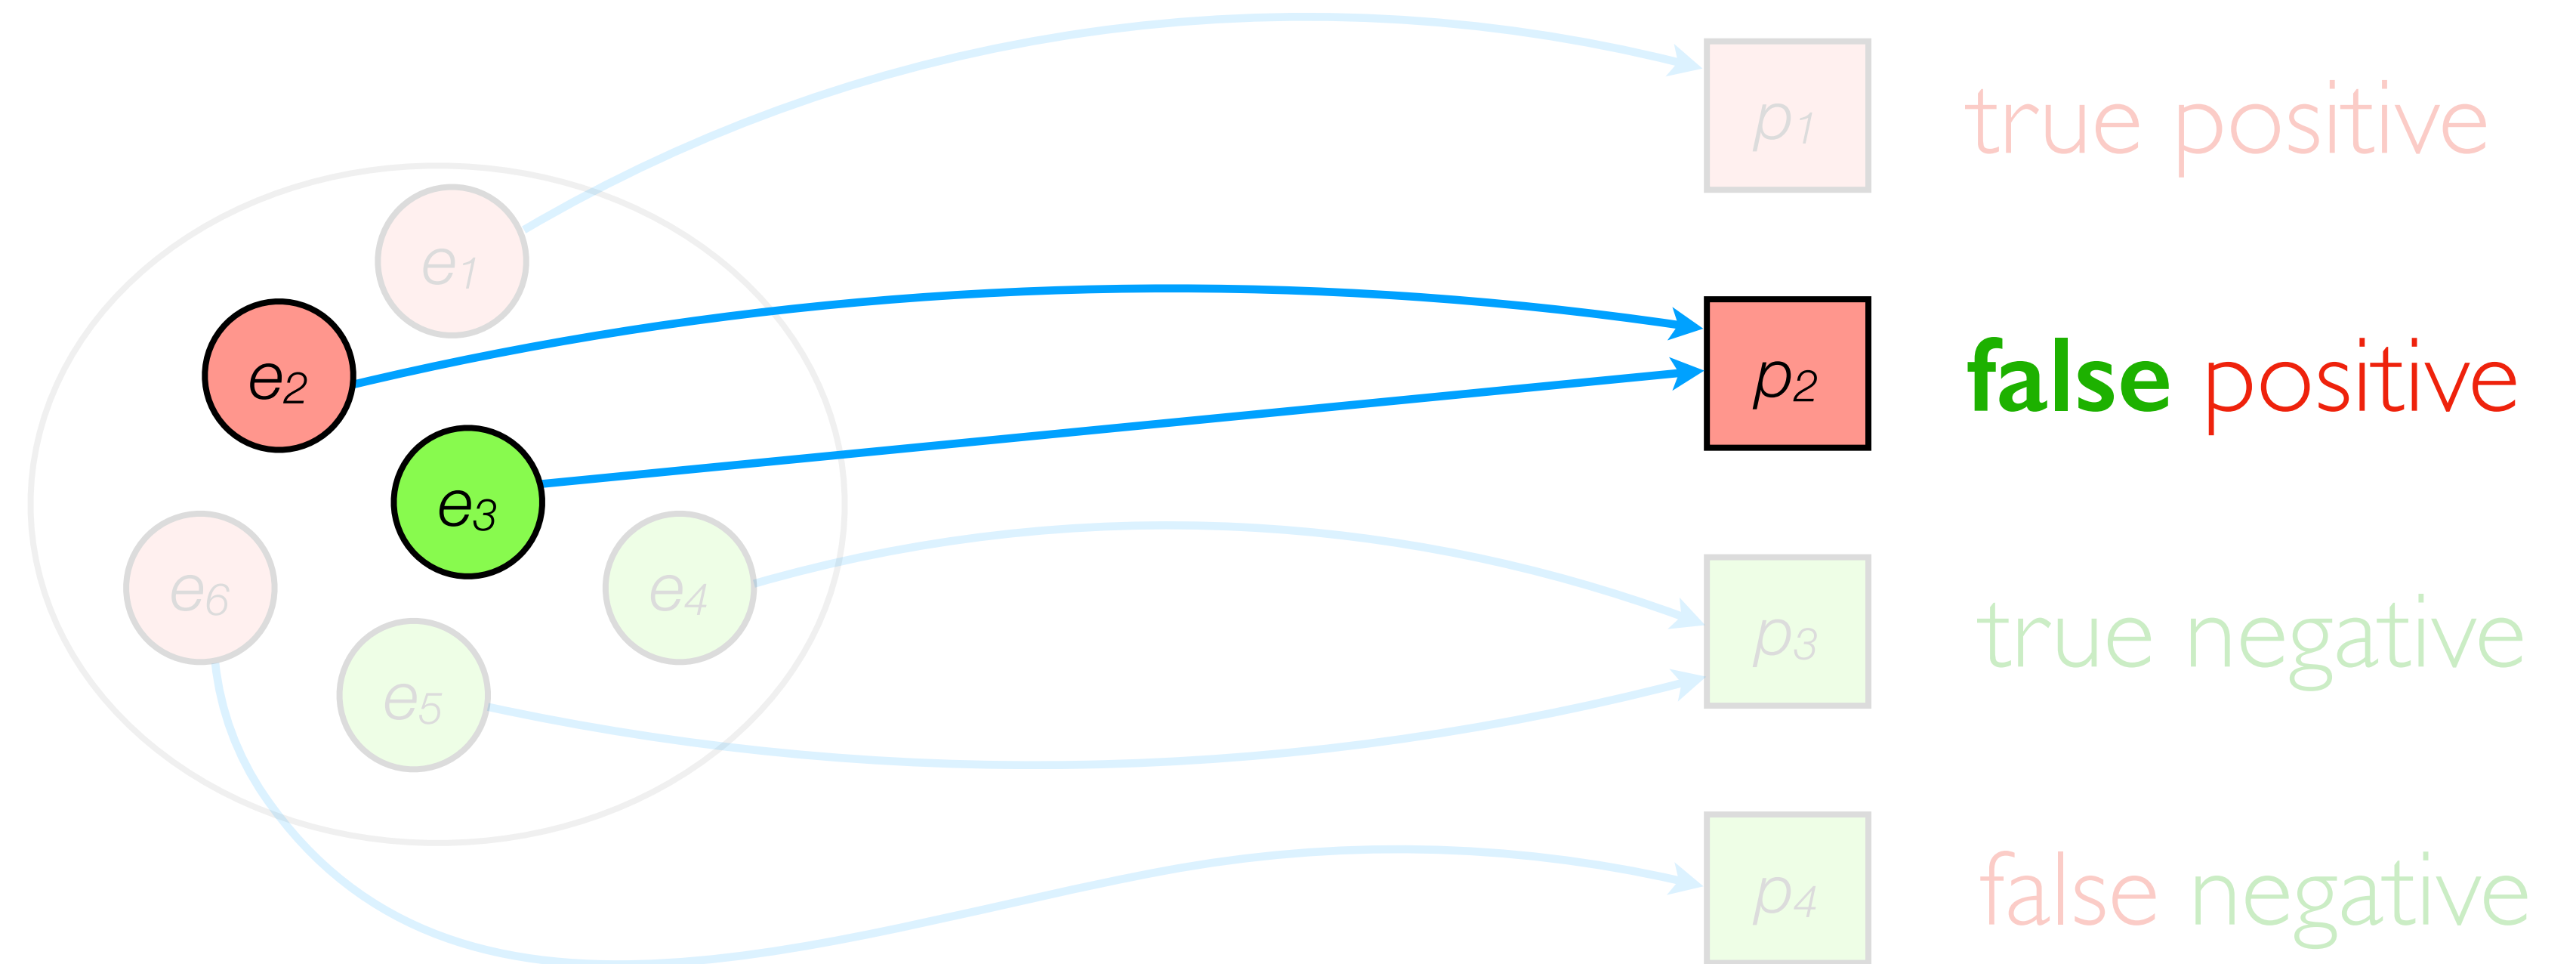
- False negatives (bugs missed) are *bad*
- False positives (non-bugs reported) are *okay*
- Constructed as ***over-approximation*** (of ***under-approximation***)
- **Soundness Theorem:**
Under certain assumptions about the programs, the analyser has *no false negatives*.



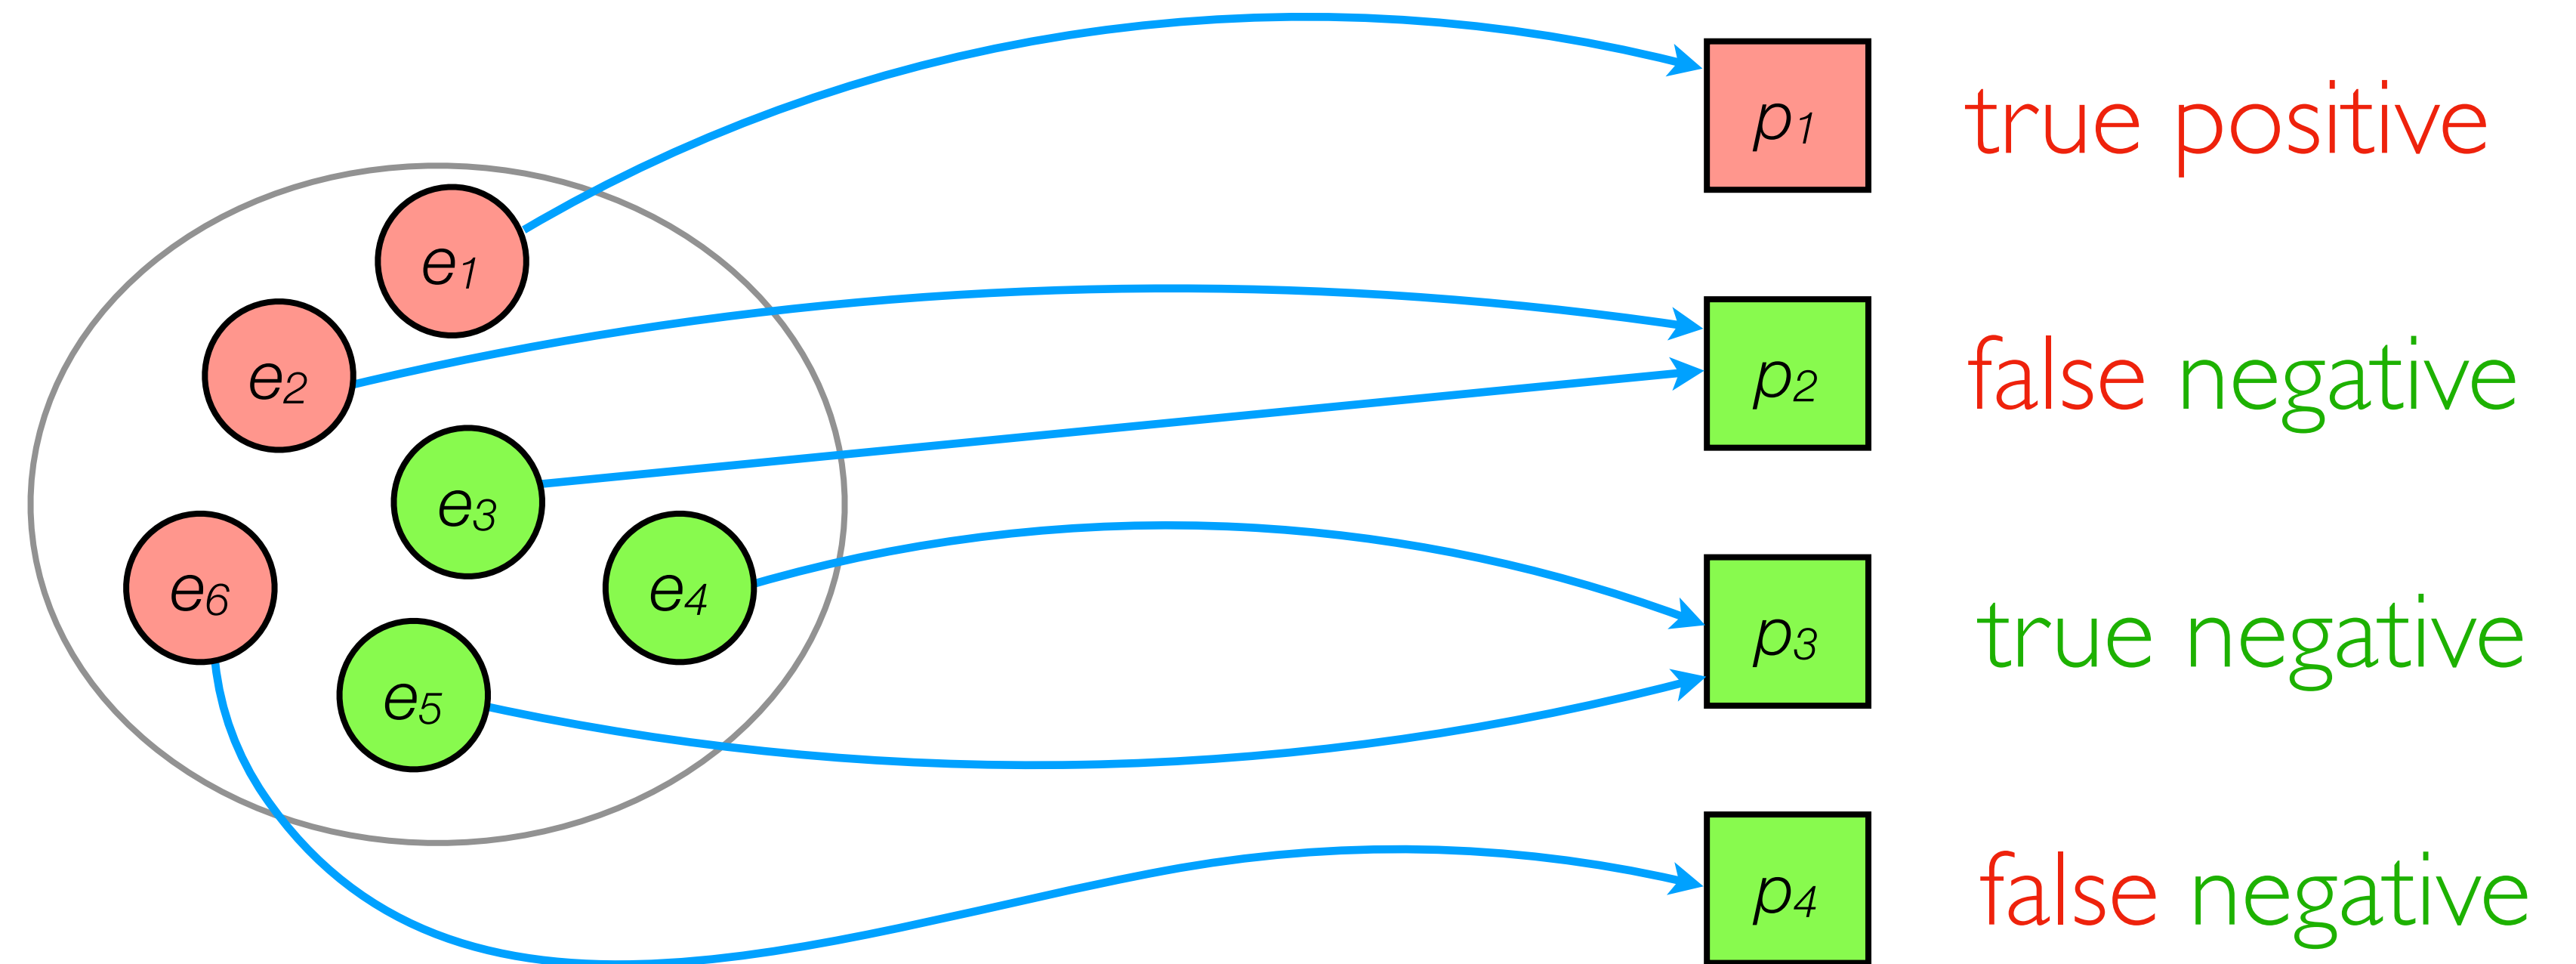
Static Bug Finder



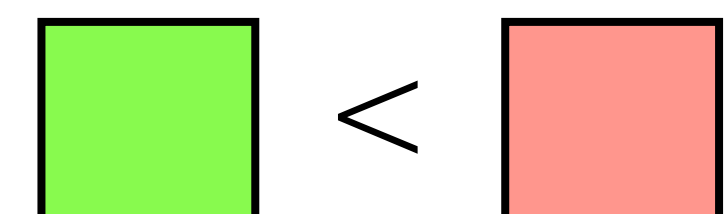
Unsound Static Bug Finder



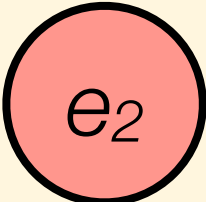
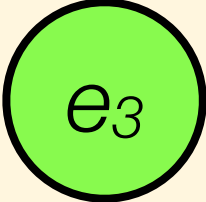
Sound (but imprecise) Static Bug Finder



abstract *under-approximation*

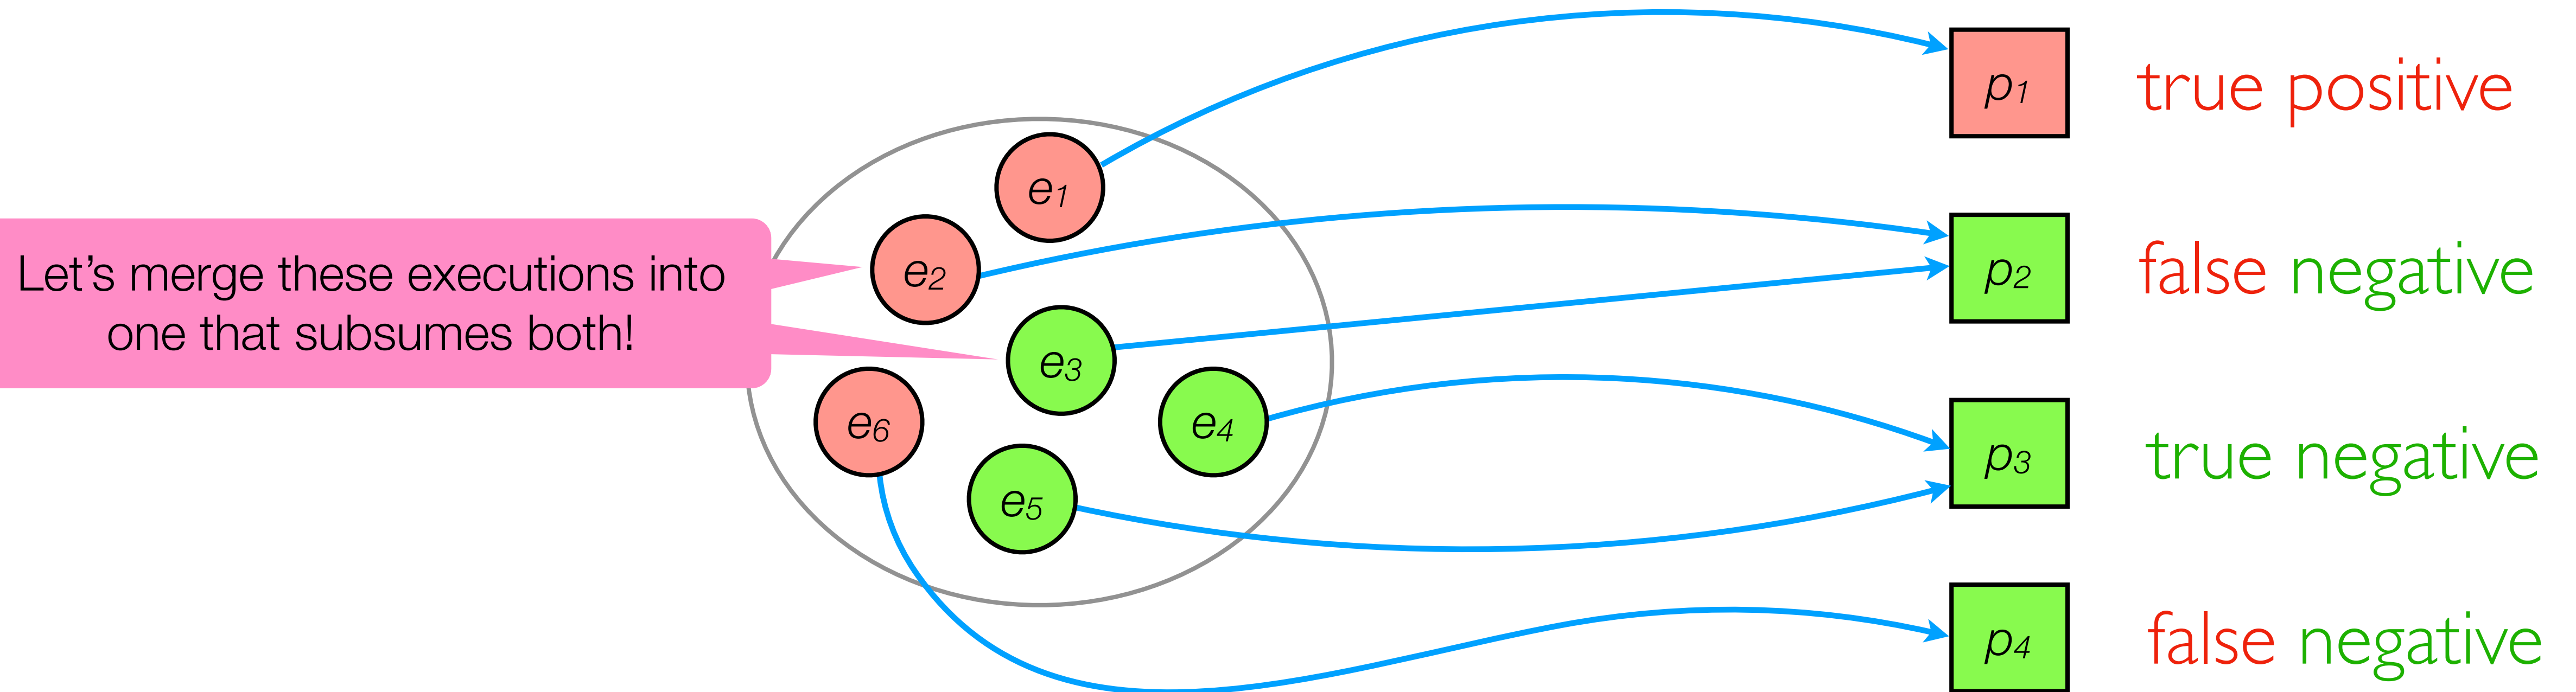


Loss of Precision in Static Bug Finders

```
if (n != VERY_UNLIKELY_VALUE) {  
     e2    // bug happens here  
} else {  
     e3    // normal execution  
}
```

Idea: *over-approximate* in concrete semantics!

Sound (but Imprecise) Static Bug Finder

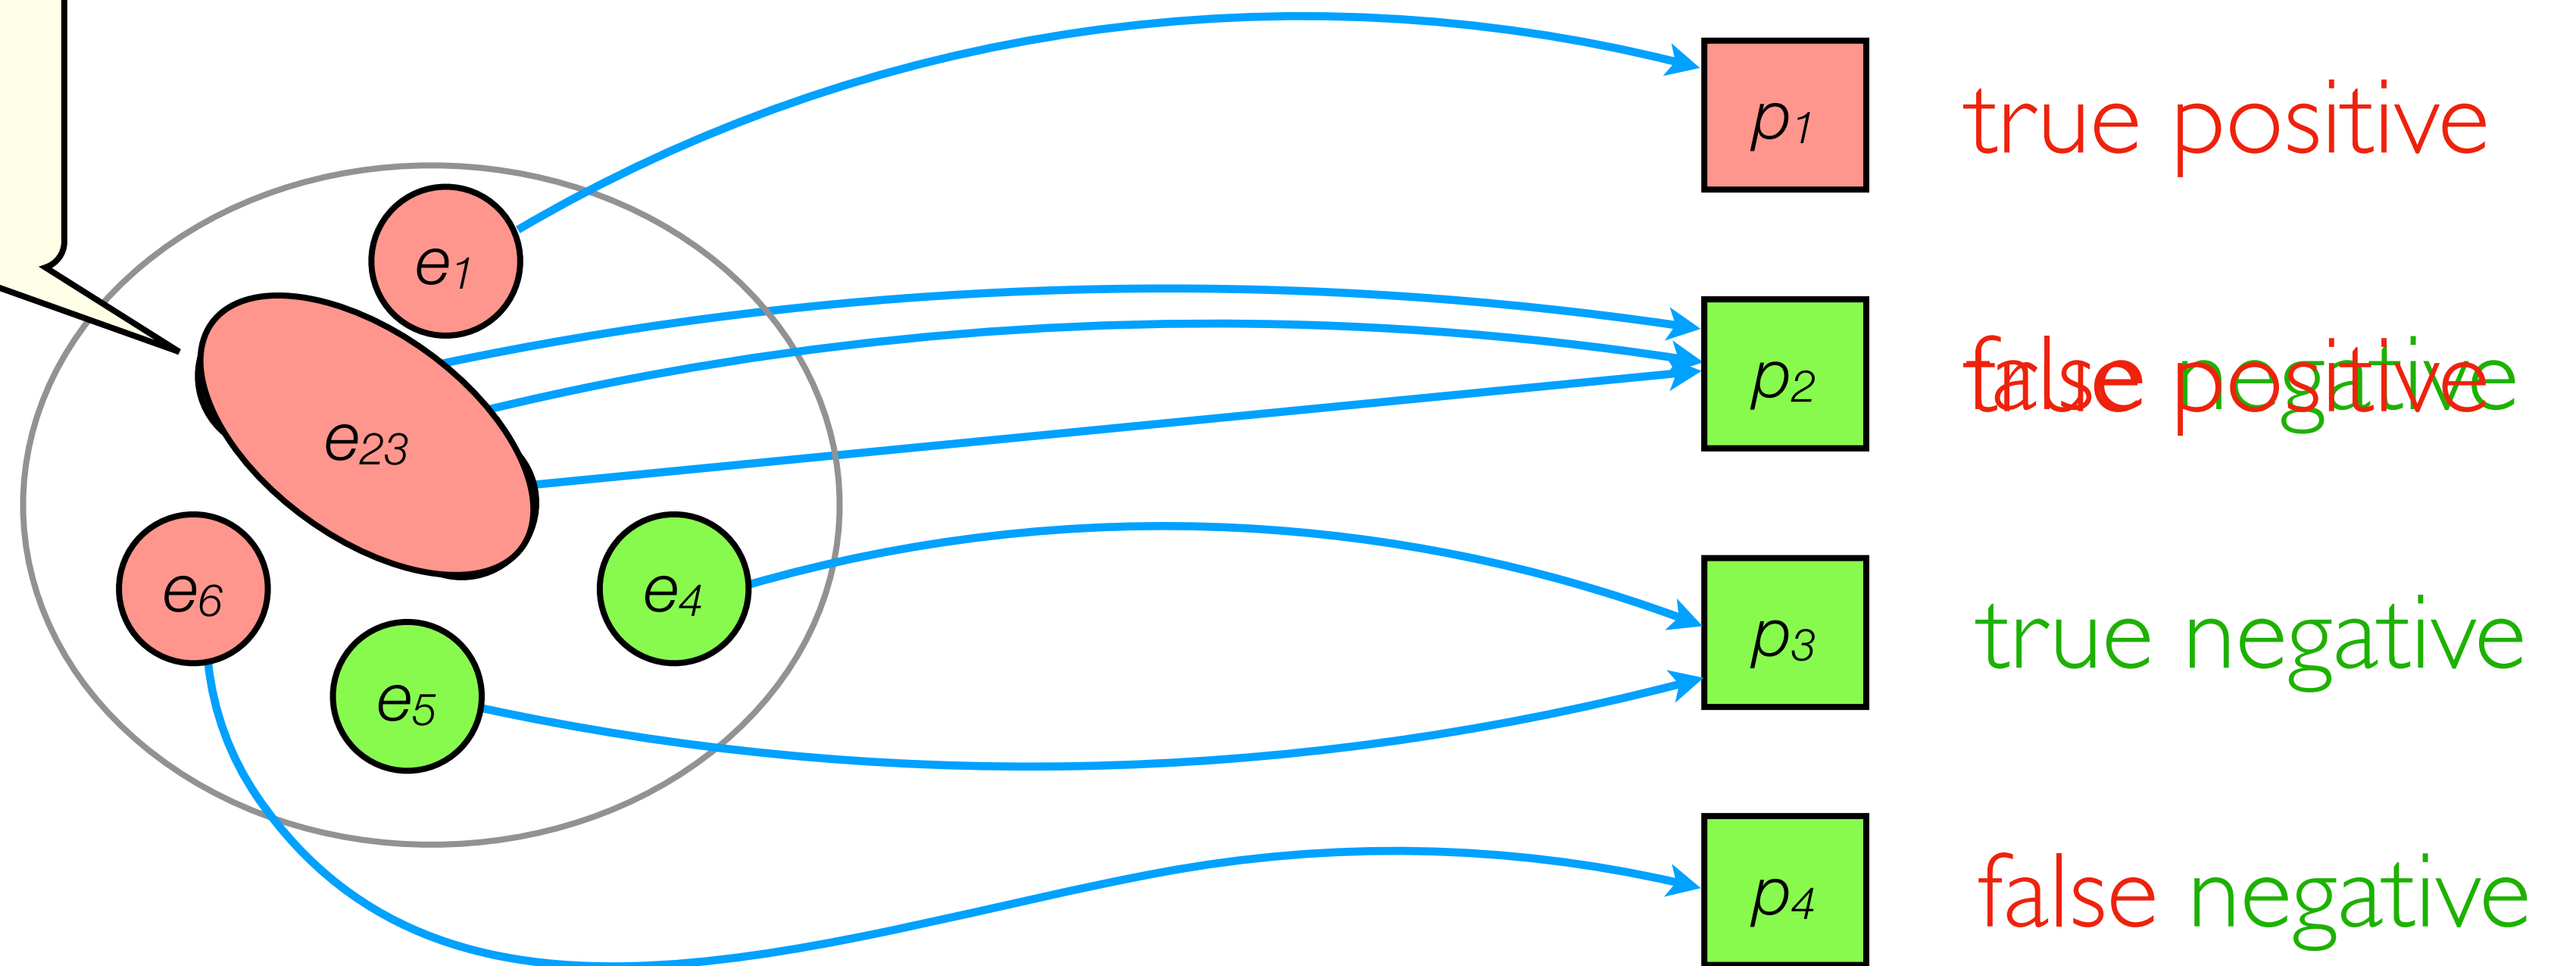


```

if (*) {
    // bug happens here
} else {
    // normal execution
}

```

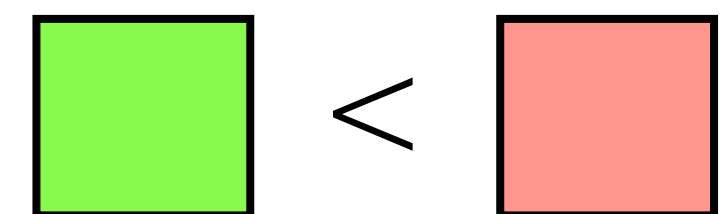
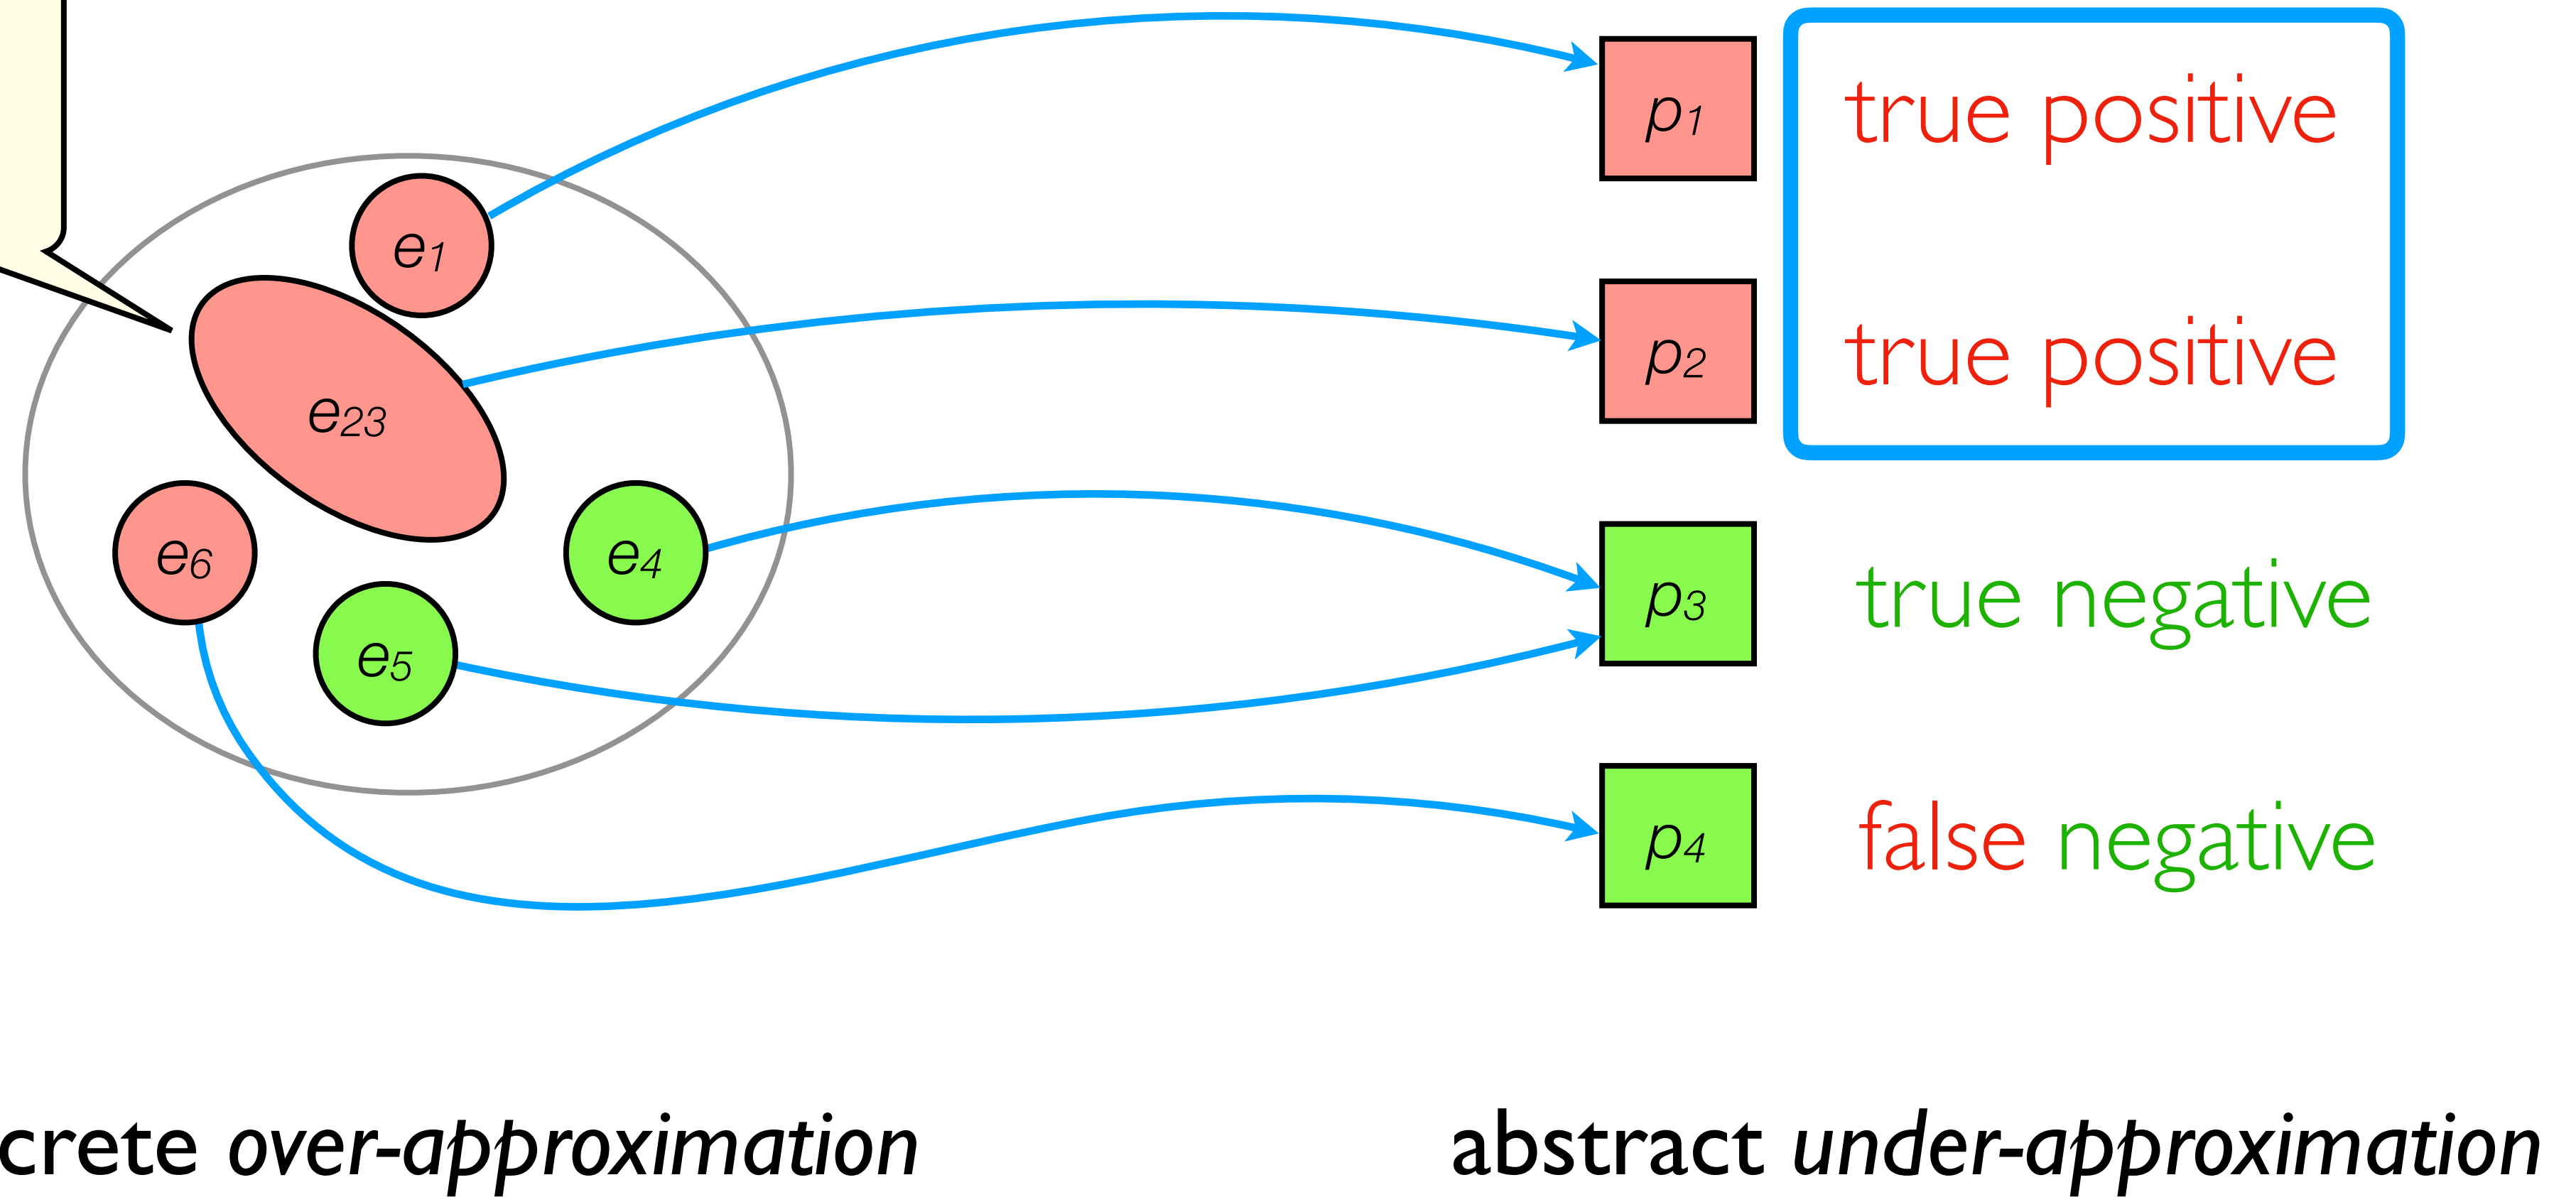
overApproxConcreteSem(c) =



Sound Static Bug Finder

```
if (*) {  
    // bug happens here  
} else {  
    // normal execution  
}
```

$\text{overApproxConcreteSem}(c) =$



Towards Sound Static Bug Finders

(this work)

- False negatives (bugs missed) are *okay*
- False positives (non-bugs reported) are *bad*
- Constructed as *under-approximation* of *over-approximation*
- Soundness (True Positives) Theorem:
Under certain assumptions about the programs, the analyser has *no false positives*.

A True Positives Theorem for a Static Race Detector



Nikos Gorogiannis

Peter O'Hearn

Ilya Sergey

facebook



facebook



YaleNUSCollege



Key Messages

Unsound (and incomplete) static analyses can be *principled*, satisfying meaningful theorems that help to understand their behaviour and guide their design

One can have an unsound but effective static analysis, which has significant industrial impact, and which is supported by a *meaningful theorem*.

Context

1. We had a demonstrably-effective industrial analysis:
RacerD (OOPSLA'18); >3k fixes in Facebook Java
2. No soundness theorem
3. Architecture: compositional abstract interpreter
4. No heuristic alarm filtering

Just ad hoc?



Our reaction:

Semantics/theory should understand/explain, not lecture.

Case Study: RacerDX

- A provably TP-Sound version of Facebook's **RacerD** concurrency analyser (Blackshear et al., OOPSLA'18)
- **Buggy executions:** *data races* in *lock-based concurrent* programs
- **Syntactic assumptions:**
Java programs with well-scoped locking (**synchronised**), no recursion, reflection, dynamic class loading; global variables are ignored.
- **Concrete over-approximation:**
Loops and conditionals are *non-deterministic*.

Formal Result

RacerDX enjoys the True Positives Theorem
wrt. Data Race Detection

(Details in the paper)



For an Idealized Language

Static Analysis with True Positives Theorem*

Goal: to build a static analysis s.t.
if the analysis reports a bug,
it is a true bug

True bug can be exhibited



The race reported
by the analysis
for program **P**
is a **true race**



There exists
an **execution** of **P**
that **exhibits**
the race



Ingredients of the formalism

For an Idealized Language

- program
- execution
- race
- analysis
- proof

Ingredients of a data race

```
class Bloop {  
    public int f = 1;  
}
```

Racy Program:

```
b = new Bloop()  
u = new Burble()
```

```
u.meps(b) || u.reps(b)
```

parallel composition

lock()
println(b.f)
unlock()

b.f := 42

```
class Burble {
```

```
    public void meps(Bloop b) {  
        synchronized (this) {  
            System.out.println(b.f);  
        }  
    }
```

path

```
    public void reps(Bloop b) {  
        b.f = 42;  
    }
```

```
    public void beps(Bloop b) {  
        b = new Bloop();  
        b.f = 239;  
    }  
}
```

Concurrent program syntax

$f \in \text{Field}$	field names
$x, \text{arg}_i \in \text{Var}$	variables
$\pi \in \text{Path} ::= x.f \mid \pi.f$	
$e \in \text{Exp} \triangleq \text{Var} \cup \text{Path}$	
$c \in \text{Stmt} ::= \text{skip} \mid x := x \mid x := \pi \mid \pi := x \mid x := \text{new}() \mid \text{lock}() \mid \text{unlock}() \mid \text{pop}()$	
$C \in \text{CStmt} ::= c \mid C; c \mid C; \text{if } * \text{ then } C \text{ else } C \mid C; \text{while } * \text{ do } C \mid C; m(e, \dots, e)$	
$M \in \text{Method} ::= m(\text{arg}_1, \dots, \text{arg}_n) \{ C \}$	
$p \in \text{Program} ::= C \parallel C$	



Single-threaded program **C**: concrete semantics

- State $\varsigma = \langle c, S, h, L \rangle$
(command, stack, heap, locks)
- Trace (list of states)
 $\tau = [\varsigma_0, \dots, \varsigma_n]$
- Concrete semantics
(set of traces)
 $\llbracket C \rrbracket \varsigma \in \mathcal{P}(\mathcal{T})$

Concrete semantics of commands

$$\llbracket \text{skip} \rrbracket \langle S, h, L \rangle \triangleq \{\epsilon\}$$

$$\llbracket x := \pi \rrbracket \langle s :: S, h, L \rangle \triangleq \begin{cases} \emptyset & \text{if } \lfloor \pi \rfloor_{s,h} \notin \text{dom}(h) \\ \{ \langle x := \pi, s[x \mapsto h(\lfloor \pi \rfloor_{s,h})] :: S, h, L \rangle \} & \text{otherwise} \end{cases}$$

$$\llbracket \pi := x \rrbracket \langle s :: S, h, L \rangle \triangleq \begin{cases} \emptyset & \text{if } \lfloor \pi \rfloor_{s,h} \notin \text{dom}(h) \\ \{ \langle \pi := x, s :: S, h[\lfloor \pi \rfloor_{s,h} \mapsto s(x)], L \rangle \} & \text{otherwise} \end{cases}$$

$$\llbracket x := \text{new}() \rrbracket \langle s :: S, h, L \rangle \triangleq \left\{ \langle x := \text{new}(), s' :: S, h', L \rangle \mid \begin{array}{l} \ell \notin \text{locn}(h), s' = s[x \mapsto \ell], \\ h' = h \cup \bigcup_{f \in \text{Field}} \{(\ell, f) \mapsto \ell\} \end{array} \right\}$$

$$\llbracket x := y \rrbracket \langle s :: S, h, L \rangle \triangleq \{ \langle x := y, s[x \mapsto s(y)] :: S, h, L \rangle \}$$

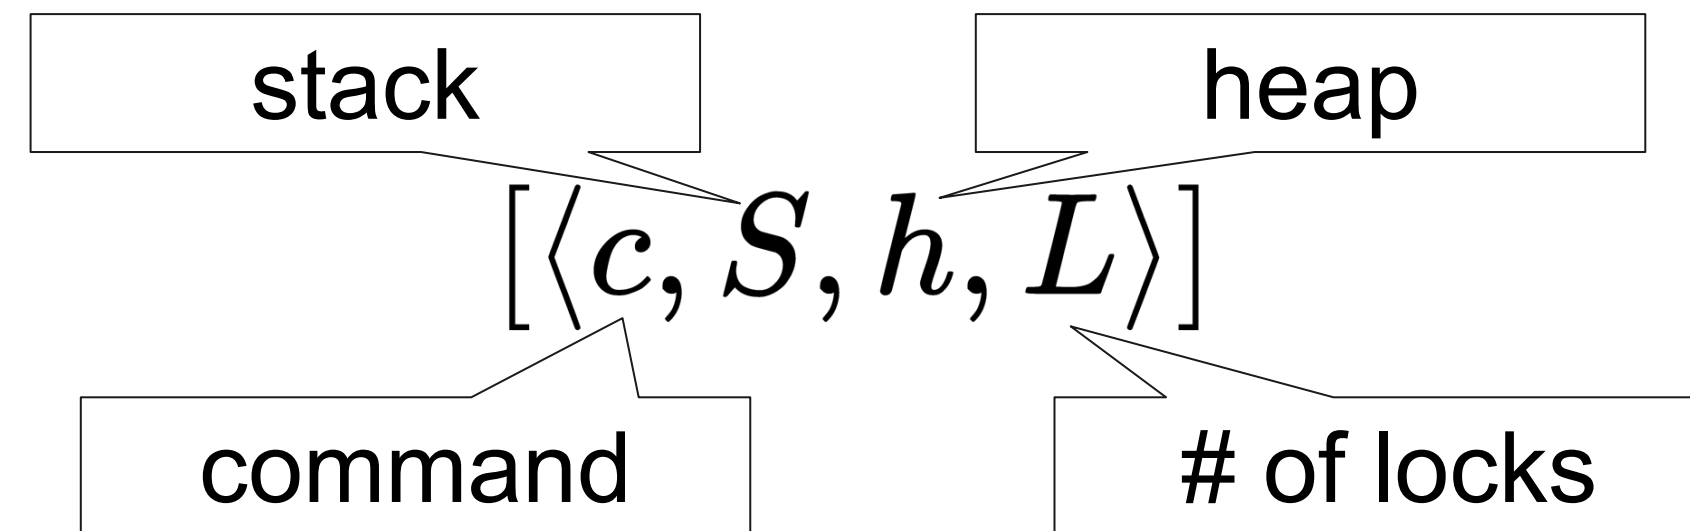
$$\llbracket \text{lock}() \rrbracket \langle S, h, L \rangle \triangleq \{ \langle \text{lock}(), S, h, \text{add}(L, 1) \rangle \}$$

$$\llbracket \text{unlock}() \rrbracket \langle S, h, L \rangle \triangleq \begin{cases} \emptyset & \text{if } L \leq 0 \\ \{ \langle \text{unlock}(), S, h, L - 1 \rangle \} & \text{otherwise} \end{cases}$$

$$\llbracket \text{pop}() \rrbracket \langle s :: S, h, L \rangle \triangleq \{ \langle \text{pop}(), S, h, L \rangle \}$$



Concrete semantics of commands



$$\llbracket \text{skip} \rrbracket \langle S, h, L \rangle \triangleq \{\epsilon\}$$

empty trace

$$\llbracket \pi := x \rrbracket \langle s :: S, h, L \rangle \triangleq \begin{cases} \emptyset & \text{if } \lfloor \pi \rfloor_{s,h} \notin \text{dom}(h) \\ \{ [\langle \pi := x, s :: S, h[\lfloor \pi \rfloor_{s,h} \mapsto s(x)], L \rangle] \} & \text{otherwise} \end{cases}$$

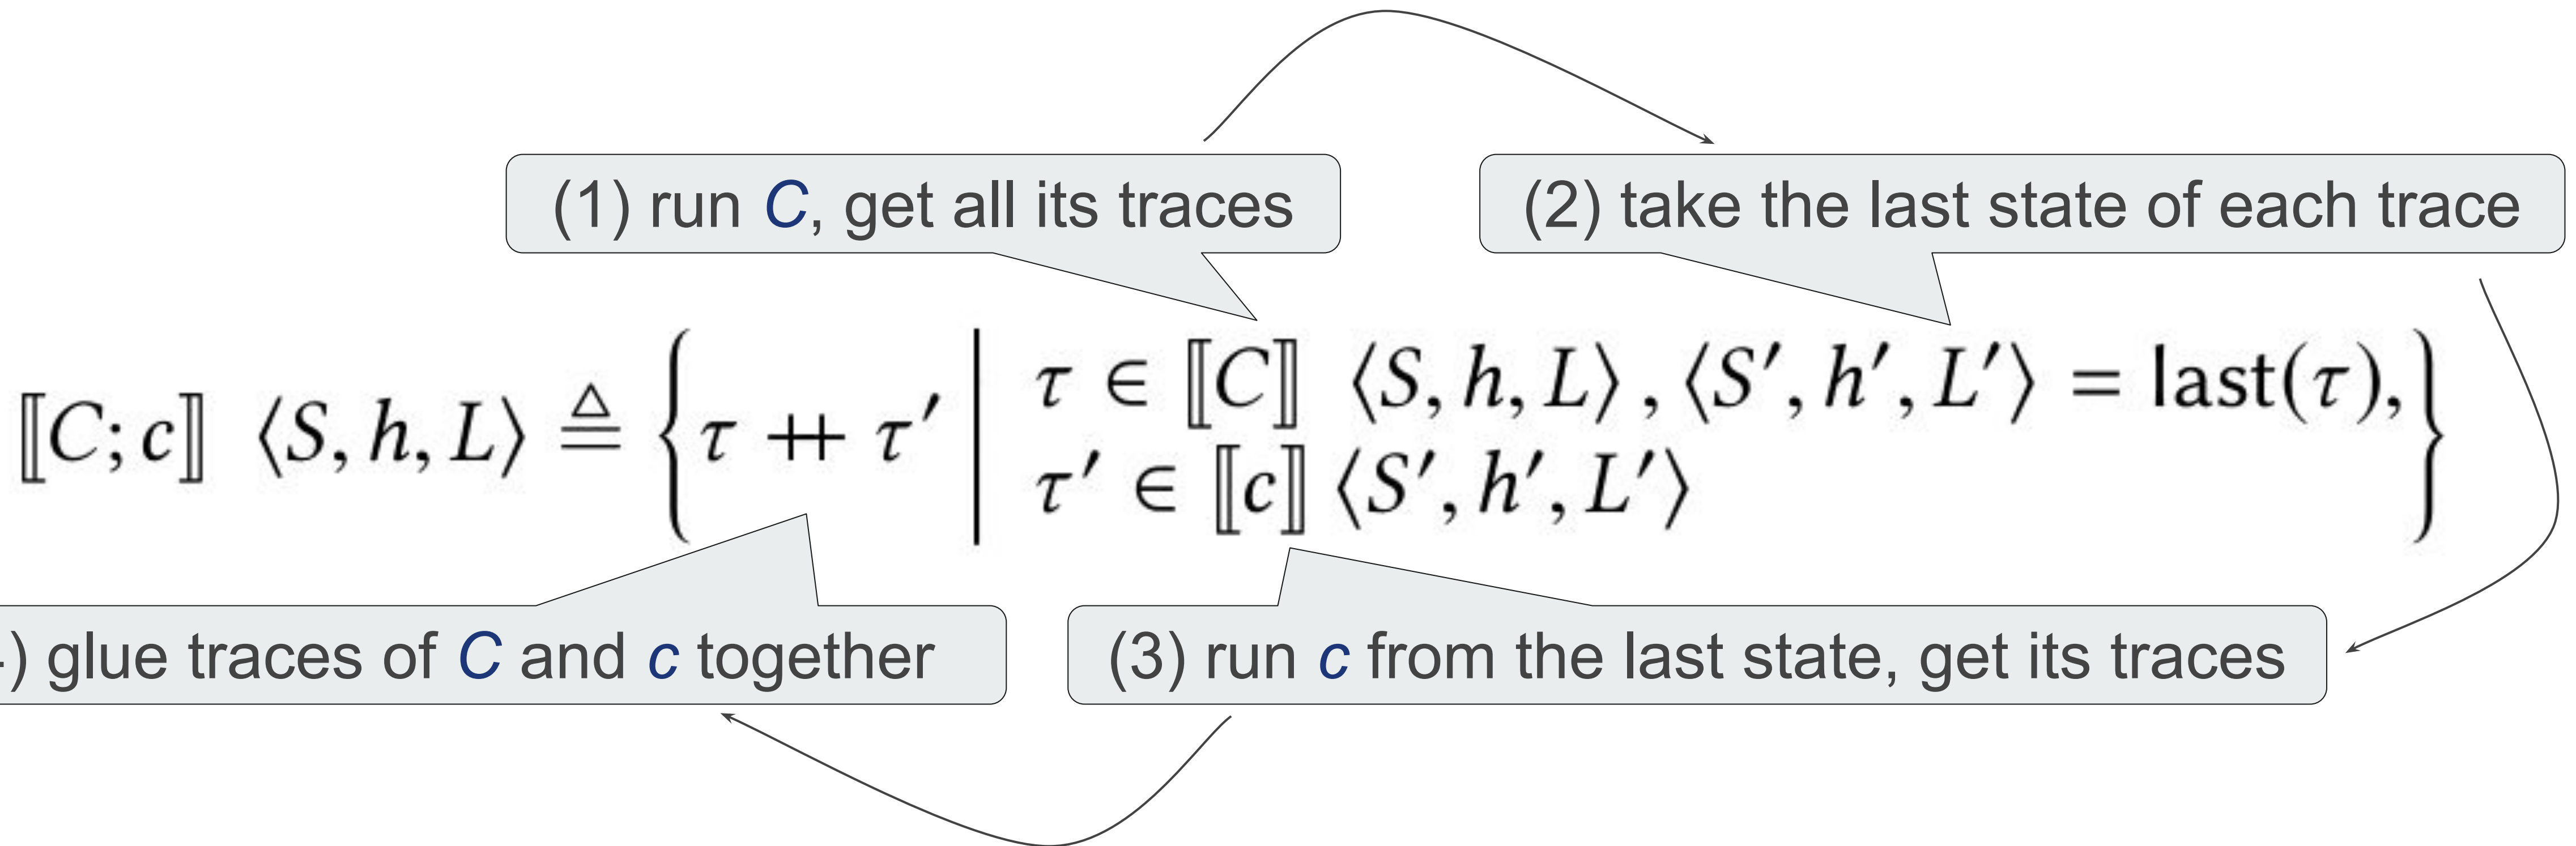
top stack frame

value of var x

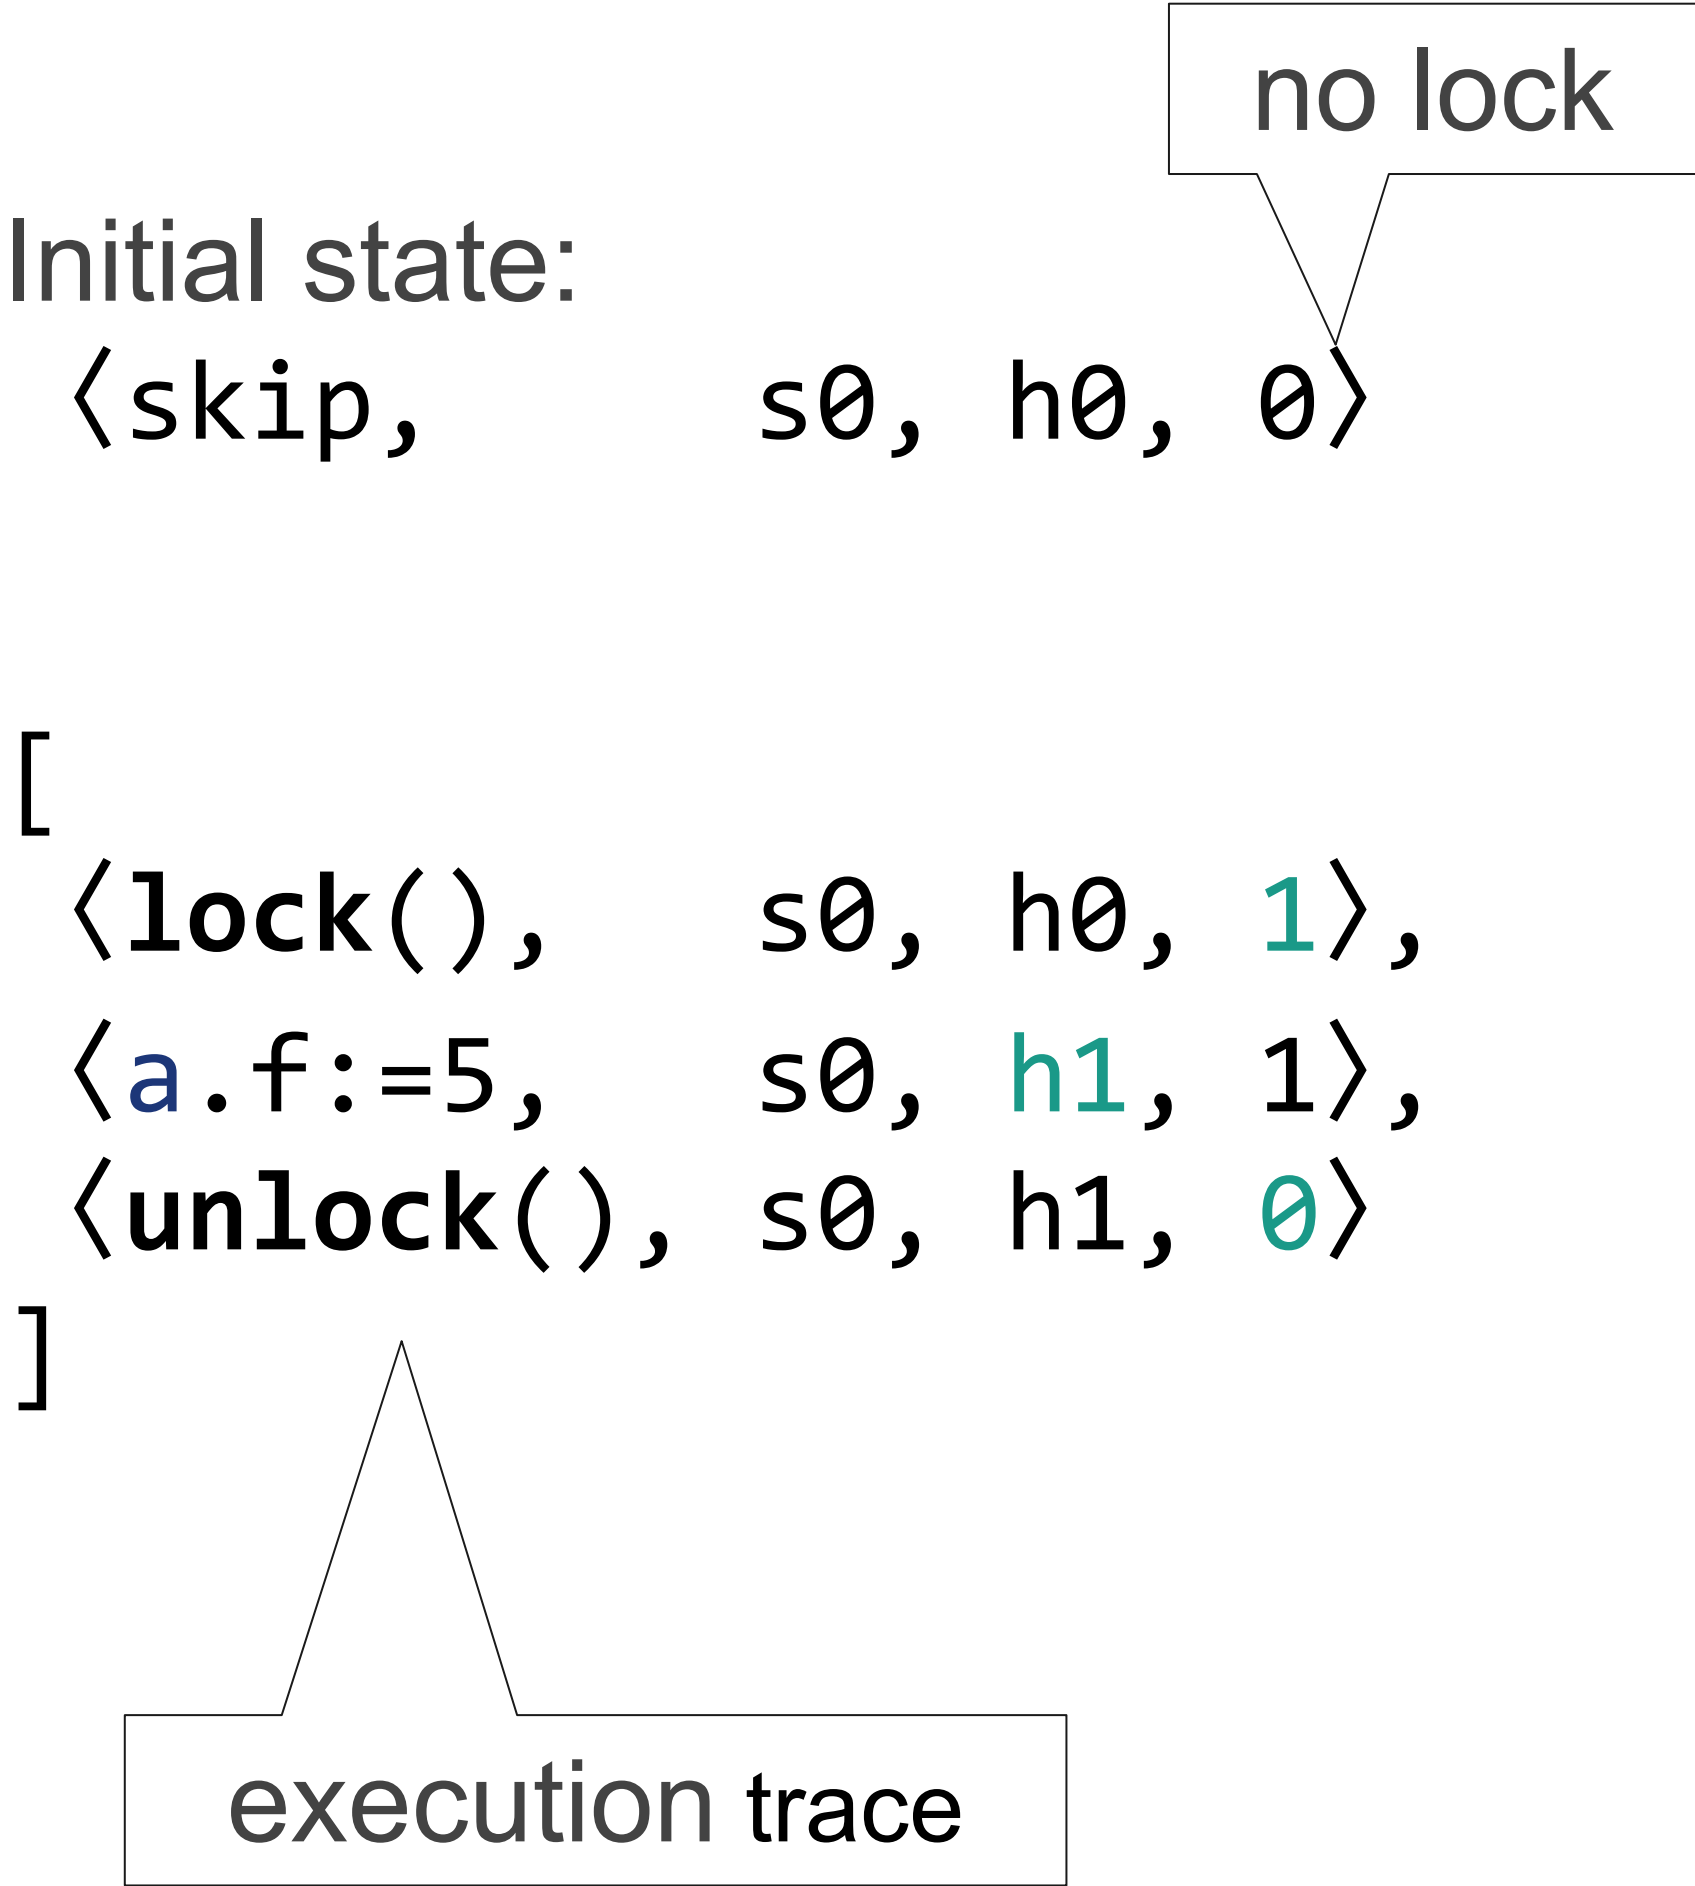
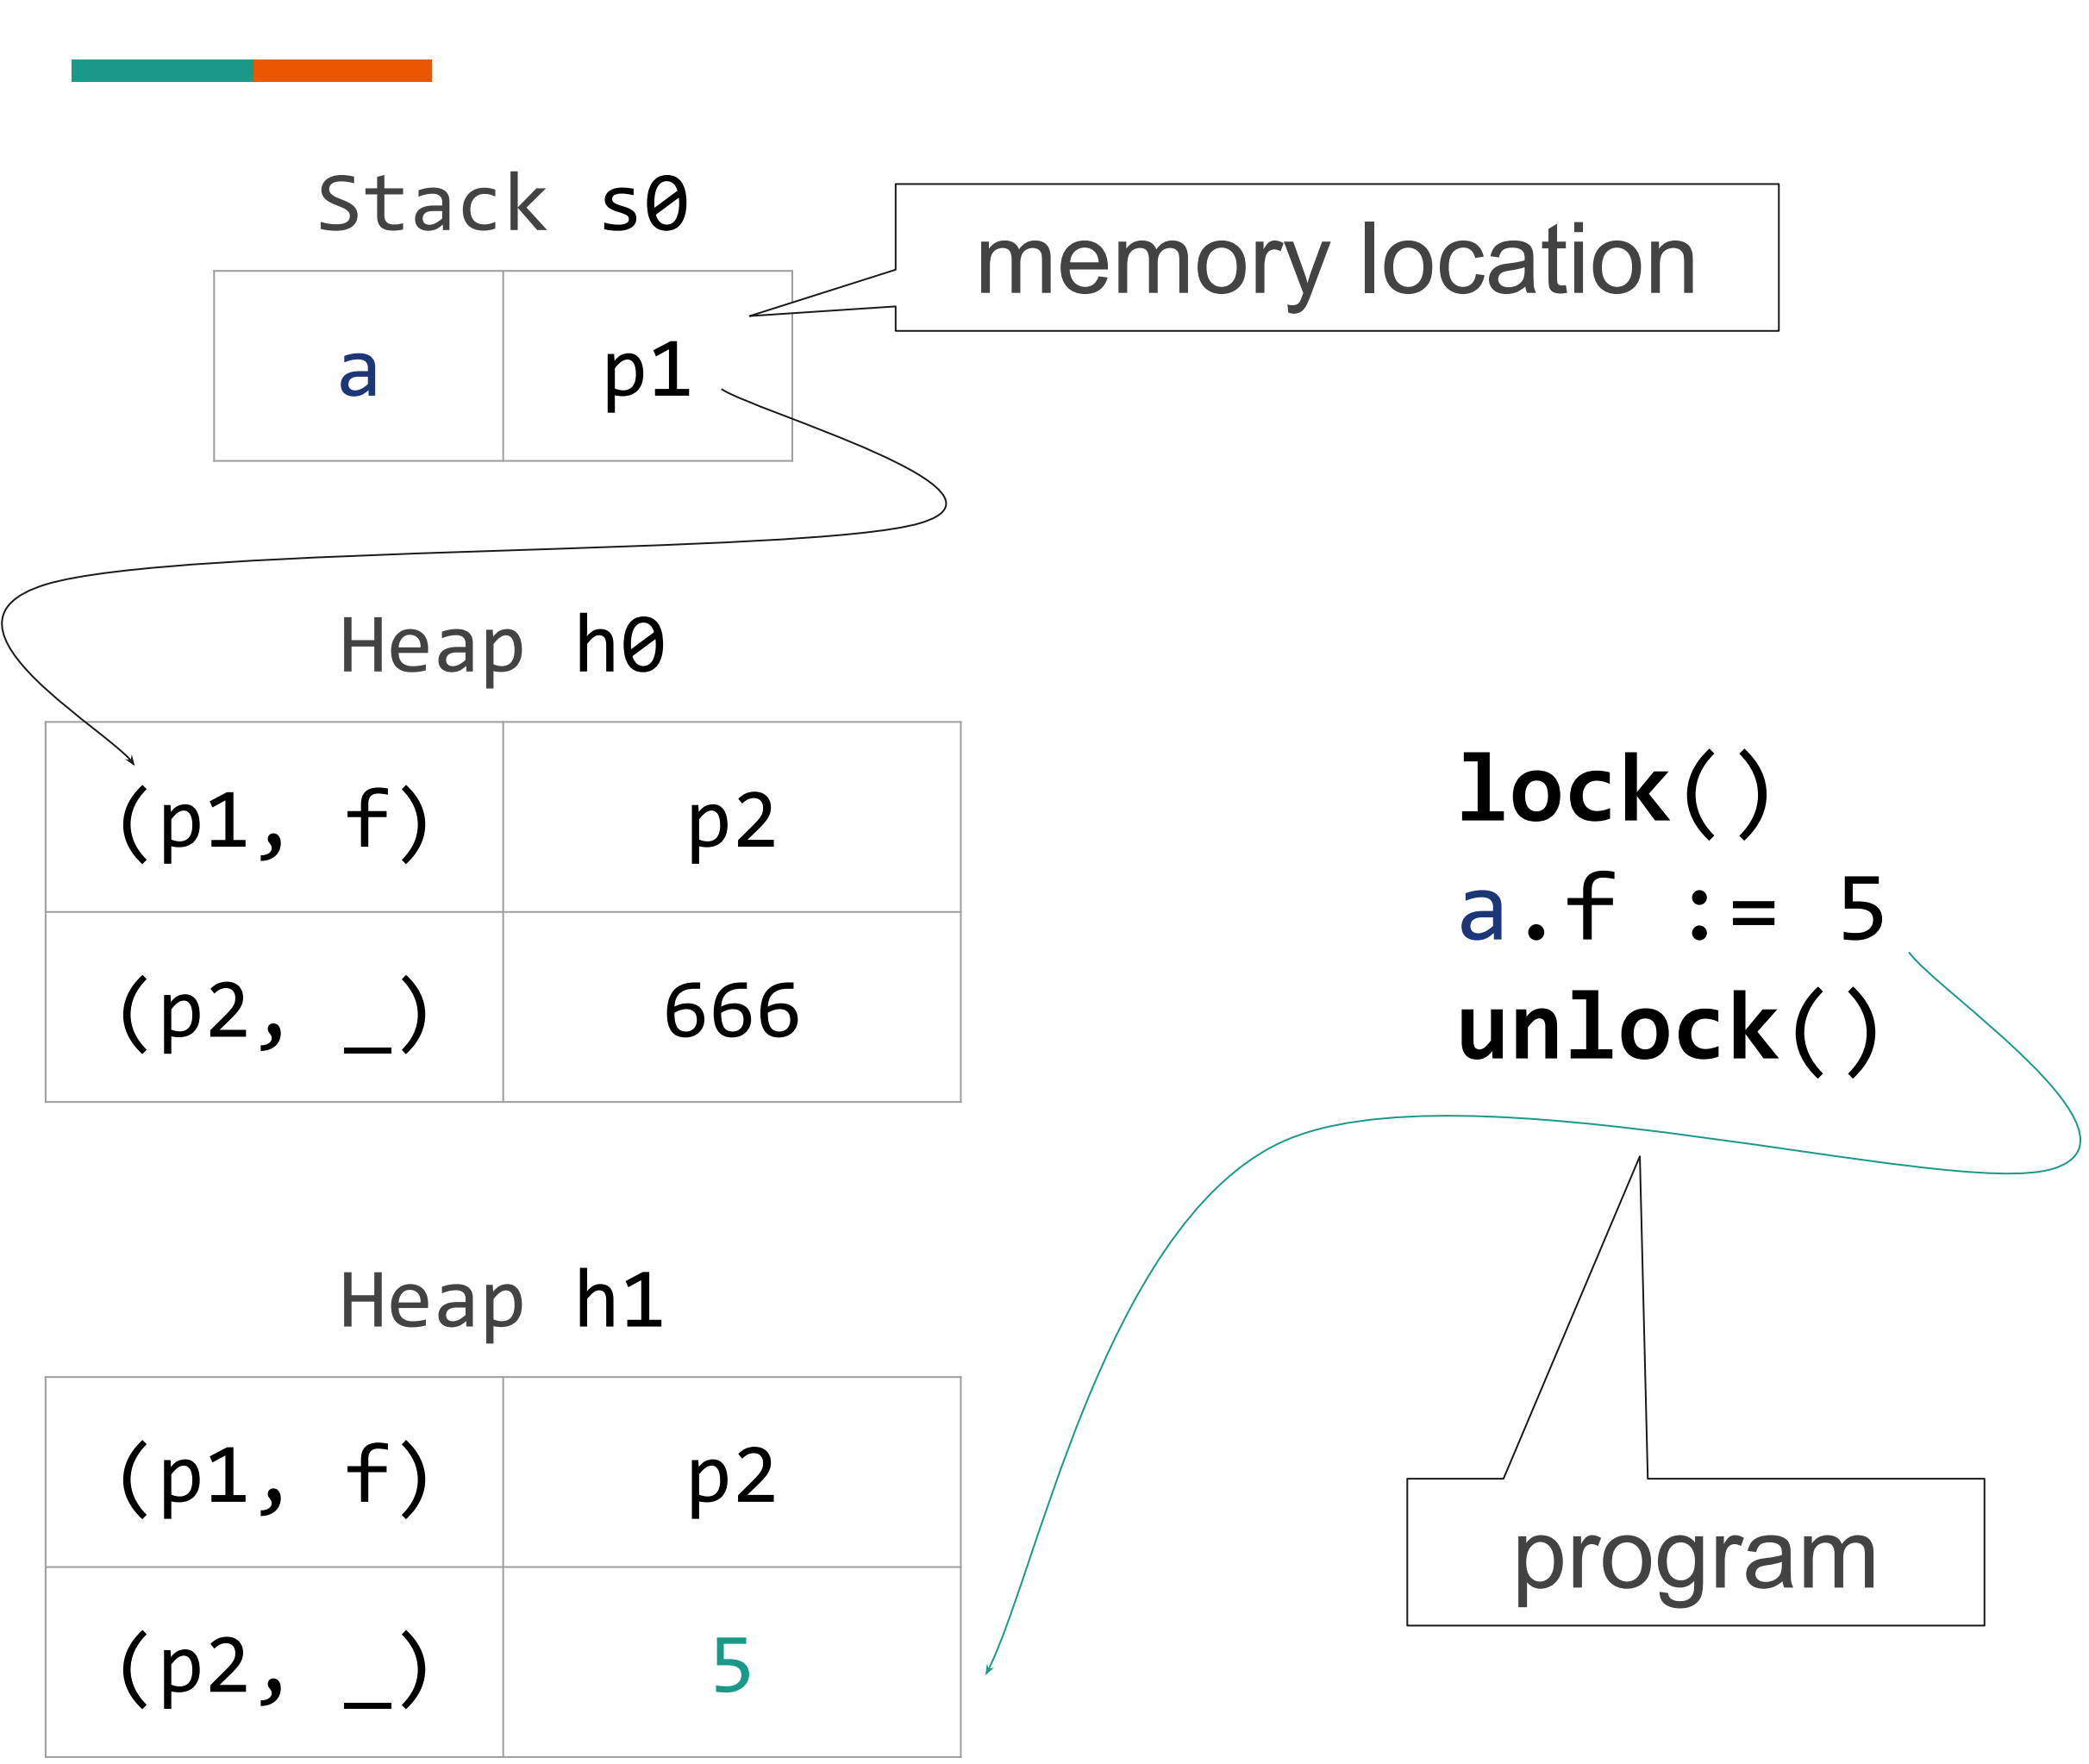
location pointed by π given stack s & heap h

$$\llbracket \text{lock}() \rrbracket \langle S, h, L \rangle \triangleq \{ [\langle \text{lock}(), S, h, \text{add}(L, 1) \rangle] \}$$

Concrete semantics of compound statements



Concrete trace example





Two-threaded program $C_1 \parallel C_2$: concrete semantics

- State

$$\langle c_{\parallel}, (S_1, S_2), h, (L_1, L_2) \rangle$$

$C \parallel \epsilon$ or $\epsilon \parallel C$

- Trace

$$\tau^{\parallel} = [\varsigma_0^{\parallel}, \dots, \varsigma_n^{\parallel}]$$

- Concrete semantics

$$\llbracket C_1 \parallel C_2 \rrbracket \varsigma^{\parallel} \in \mathcal{P}(\mathcal{T}^{\parallel})$$

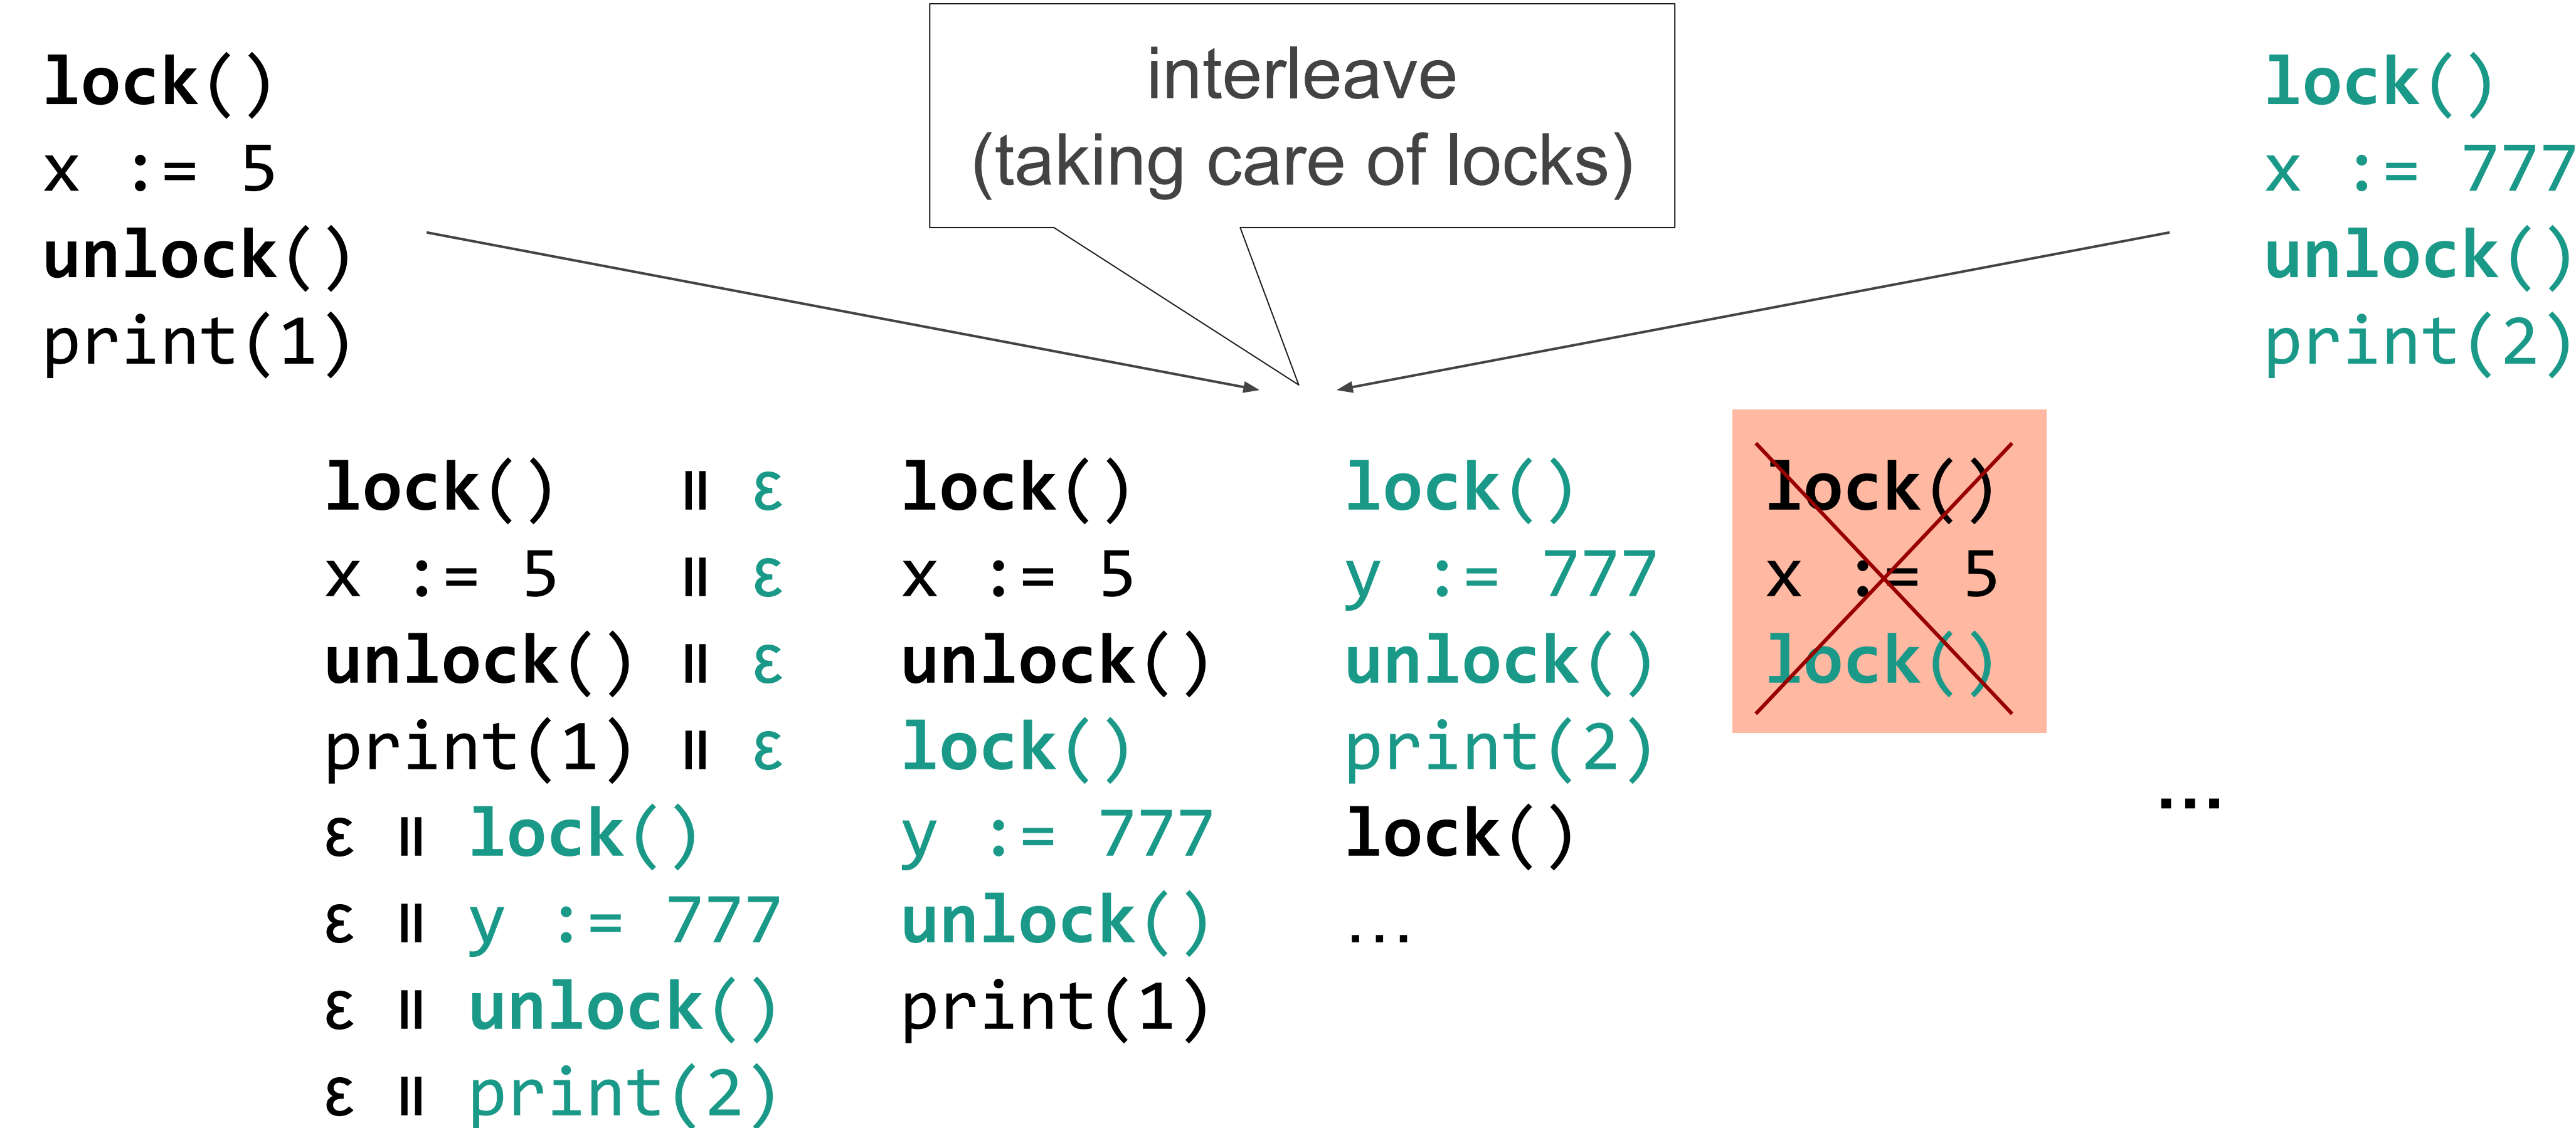
2-threaded program interleaves single traces

$$\begin{aligned} & \llbracket C_1 \parallel C_2 \rrbracket \langle (S_1, S_2), h, (L_1, L_2) \rangle \\ & \quad \triangleq \\ & \bigcup_{\substack{\tau_1 \in \llbracket C_1 \rrbracket \langle S_1, h, L_1 \rangle \\ \tau_2 \in \llbracket C_2 \rrbracket \langle S_2, h, L_2 \rangle}} \text{interl } \tau_1 \ \tau_2 \ \langle (S_1, S_2), h, (L_1, L_2) \rangle \end{aligned}$$

(1) run components individually

(2) interleave all individual traces (full and partial)

Concurrent traces example



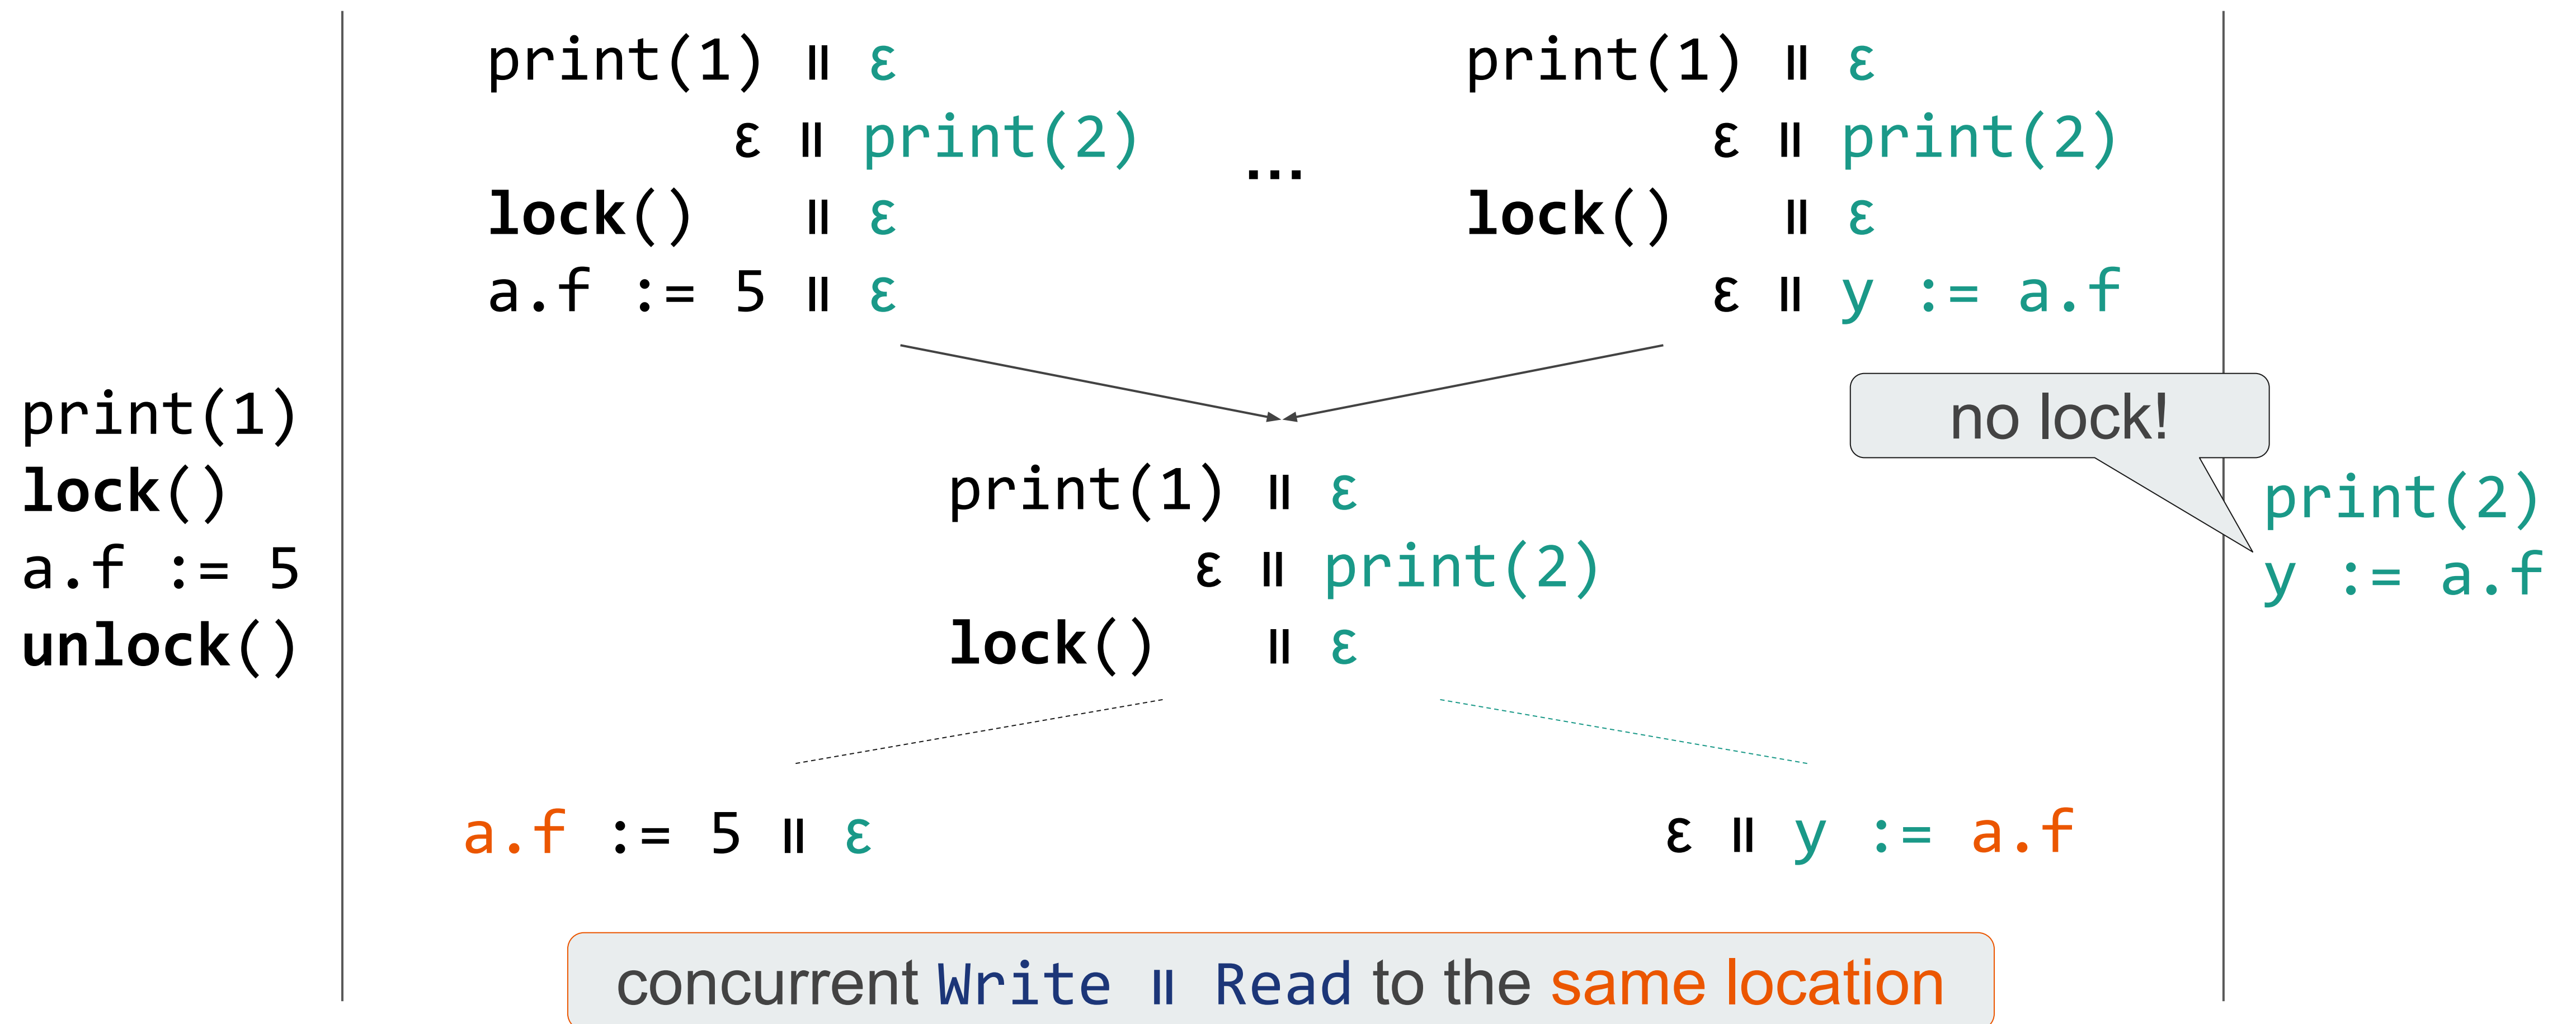
Data race means concurrent access to location

Definition 3.6 (Data Race). The program $C_1 \parallel C_2$ *races* if there exists a state ς_0 and a non-empty concurrent trace $\tau \in \llbracket C_1 \parallel C_2 \rrbracket \varsigma_0$ such that $\text{last}(\tau) = \langle _, (s_1 \mathbin{::} _, s_2 \mathbin{::} _), h, _ \rangle^9$ and,

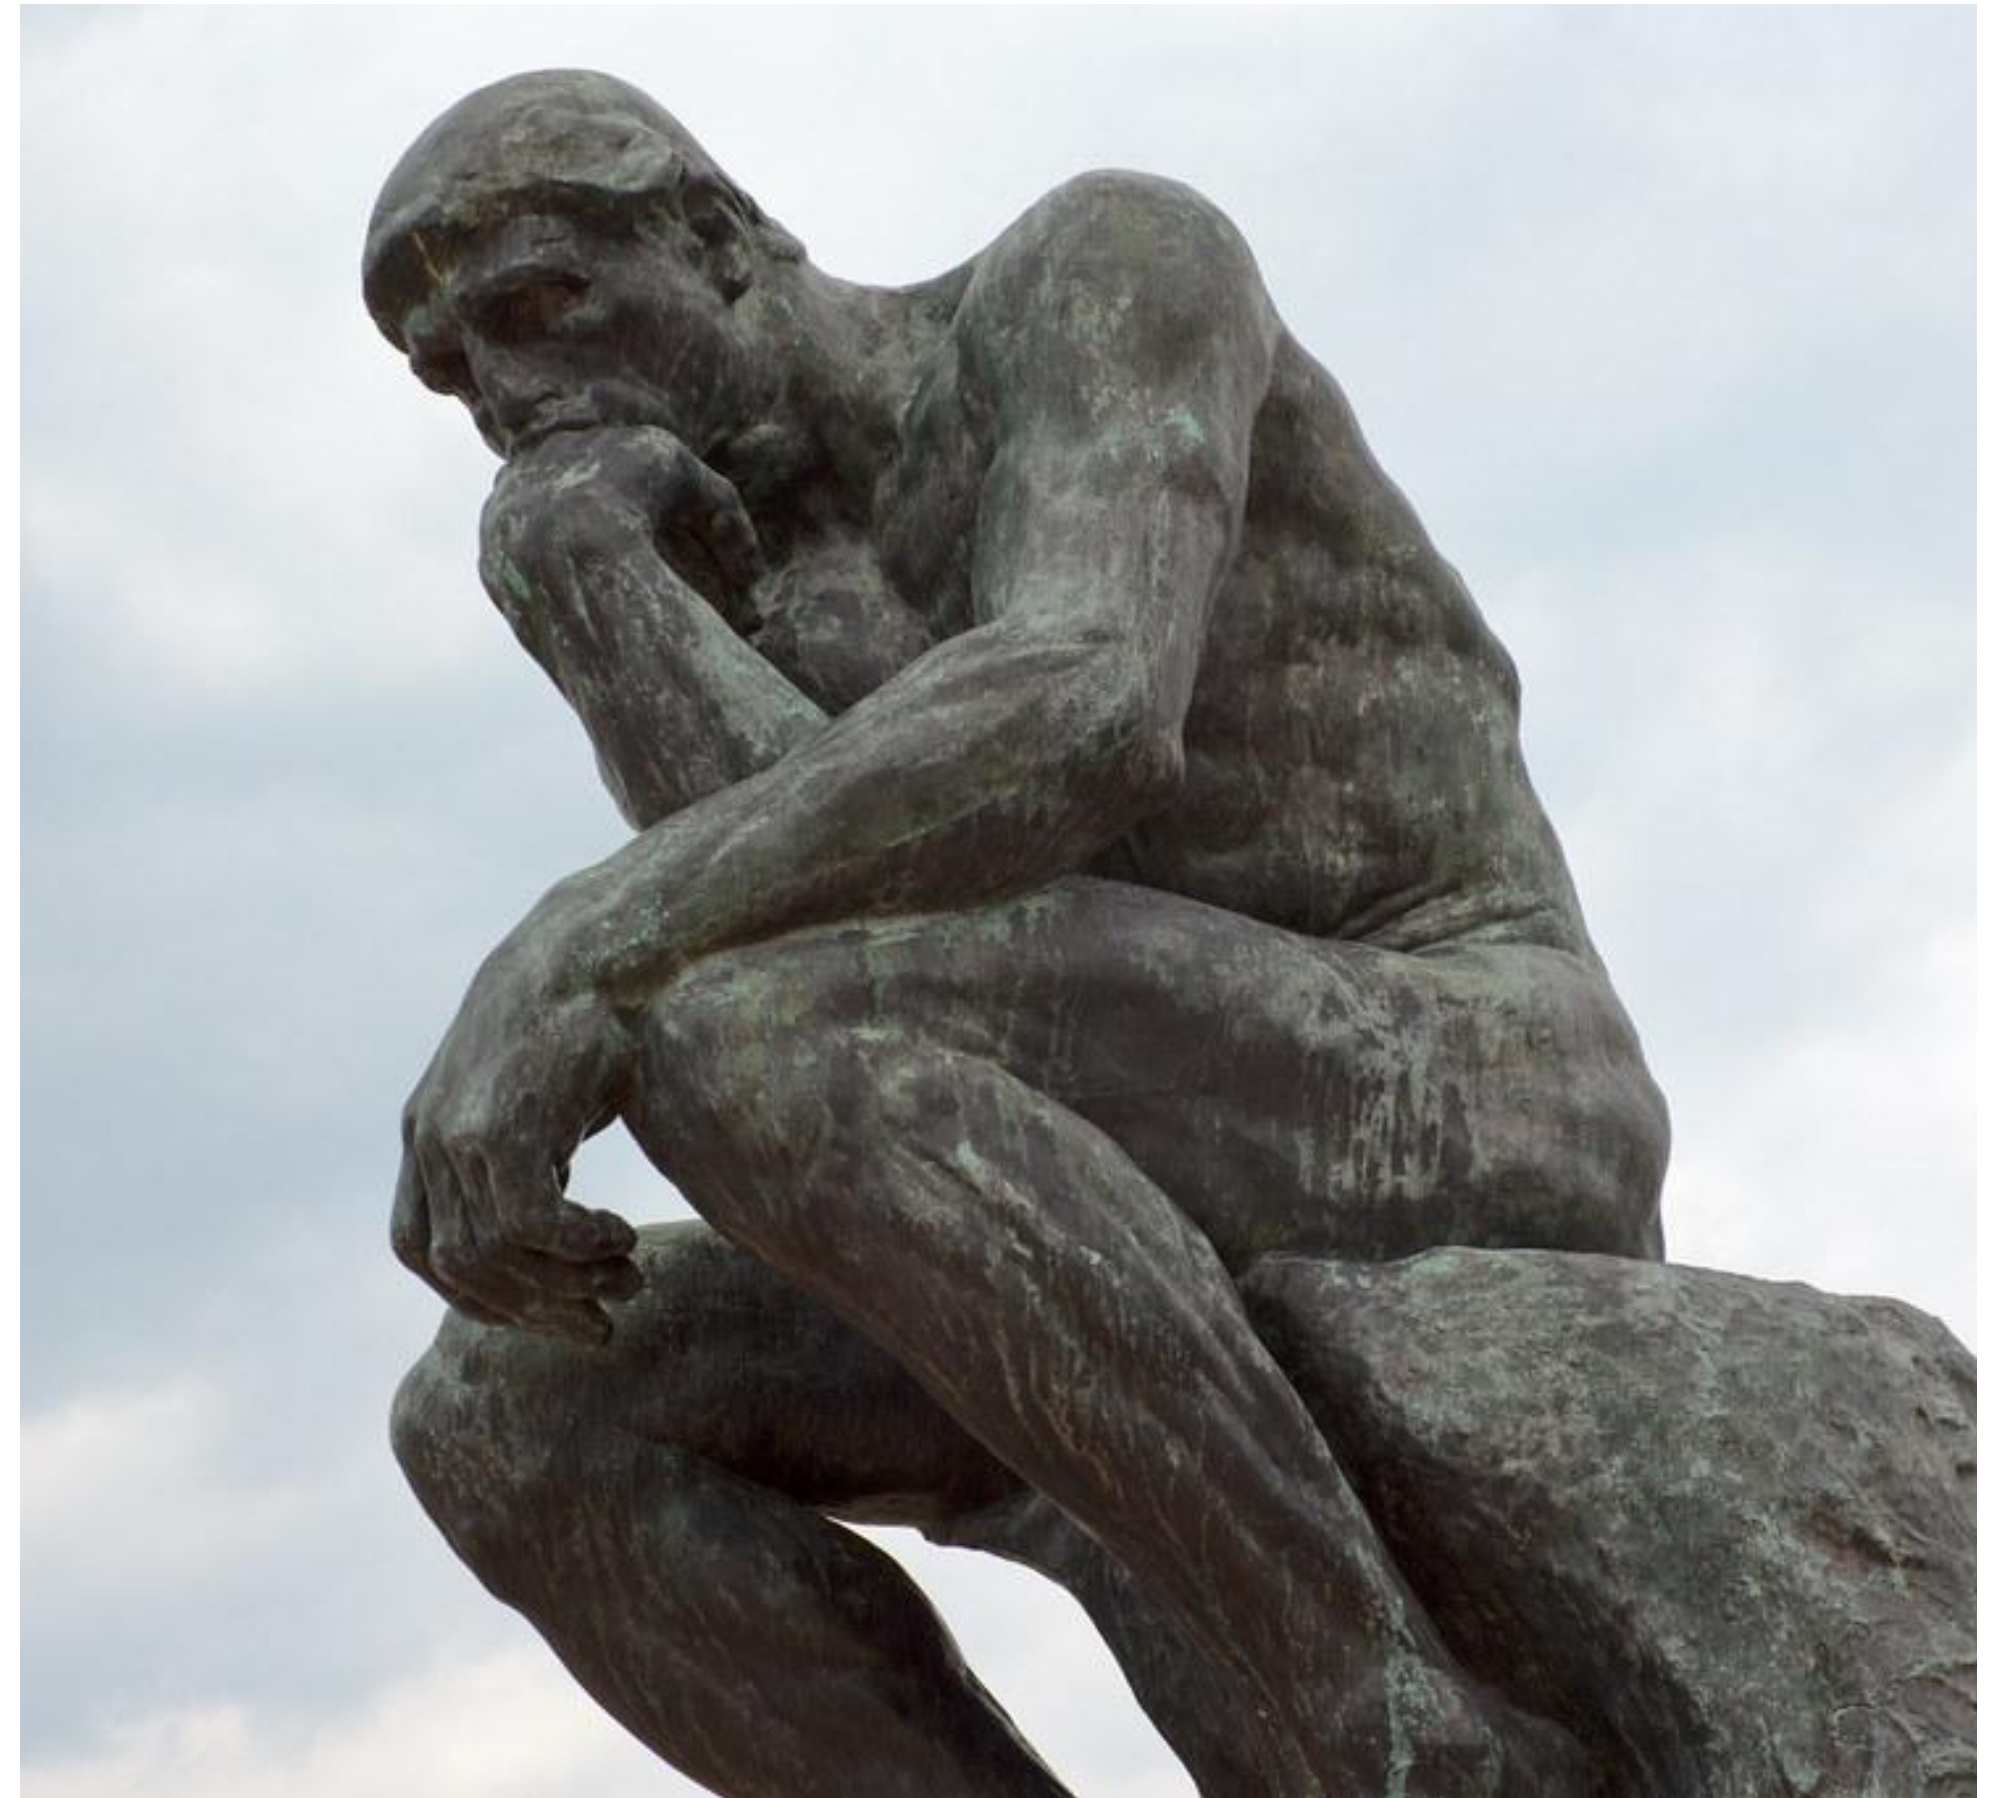
- there exist paths π_1, π_2 such that $\lfloor \pi_1 \rfloor_{s_1, h} = \lfloor \pi_2 \rfloor_{s_2, h}$;
- there exist states $\varsigma_1 = \langle c_1 \parallel \epsilon, _, _, _ \rangle$ and $\varsigma_2 = \langle \epsilon \parallel c_2, _, _, _ \rangle$ such that $\tau \mathbin{::} \varsigma_1, \tau \mathbin{::} \varsigma_2 \in \llbracket C_1 \parallel C_2 \rrbracket \varsigma_0$;
- $c_1 = (\pi_1 := _) \wedge c_2 = (\pi_2 := _)$, or, $c_1 = (\pi_1 := _) \wedge c_2 = (_ := \pi_2)$, or, $c_1 = (_ := \pi_1) \wedge c_2 = (\pi_2 := _)$.



Data race means concurrent access to location



**Can we identify
a data race
without building
the traces?**





Abstract Semantics

tracks accesses
to memory locations

helps identify
true races

- Abstract State $\langle W, L, A \rangle$
(wobblies, locks, path accesses)
- Abstraction of a set of
concrete *single*-threaded traces

$$\alpha(T) = \langle W, L, A \rangle$$

Wobblies can evade data races (produce false positives)

```
class Bloop {  
    public int f = 1;  
}
```

same path `b.f` refers to **different locations**

Path prefix `b` is “*unstable*” (“*wobbly*”),
as it’s reassigned, hence race is evaded.

```
class Burble {  
  
    public void meps(Bloop b) {  
        synchronized (this) {  
            System.out.println(b.f);  
        }  
    }  
  
    public void reps(Bloop b) {  
        b.f = 42;  
    }  
  
    public void beps(Bloop b) {  
        b = new Bloop();  
        b.f = 239;  
    }  
}
```


Abstraction keeps track of accesses and wobbles

τ without heap, stack, locks

substitution (needed for method calls)

$\text{exec } \hat{\tau} (\langle W, L, A \rangle, \bar{\theta})$

(3) mark x & π as wobbly

(2) remember access to π

$\text{exec } \hat{\tau} :: \langle \pi := x \rangle (\langle W, L, A \rangle, \bar{\theta}) \triangleq \left(\left\langle W' \cup \{x, \pi\}^{\bar{\theta}'}, L', A' \cup \{ \langle \pi := x, L' \rangle \}^{\bar{\theta}'} \right\rangle, \bar{\theta}' \right)$

where $(\langle W', L', A' \rangle, \bar{\theta}') \triangleq \text{exec } \hat{\tau} (\langle W, L, A \rangle, \bar{\theta})$

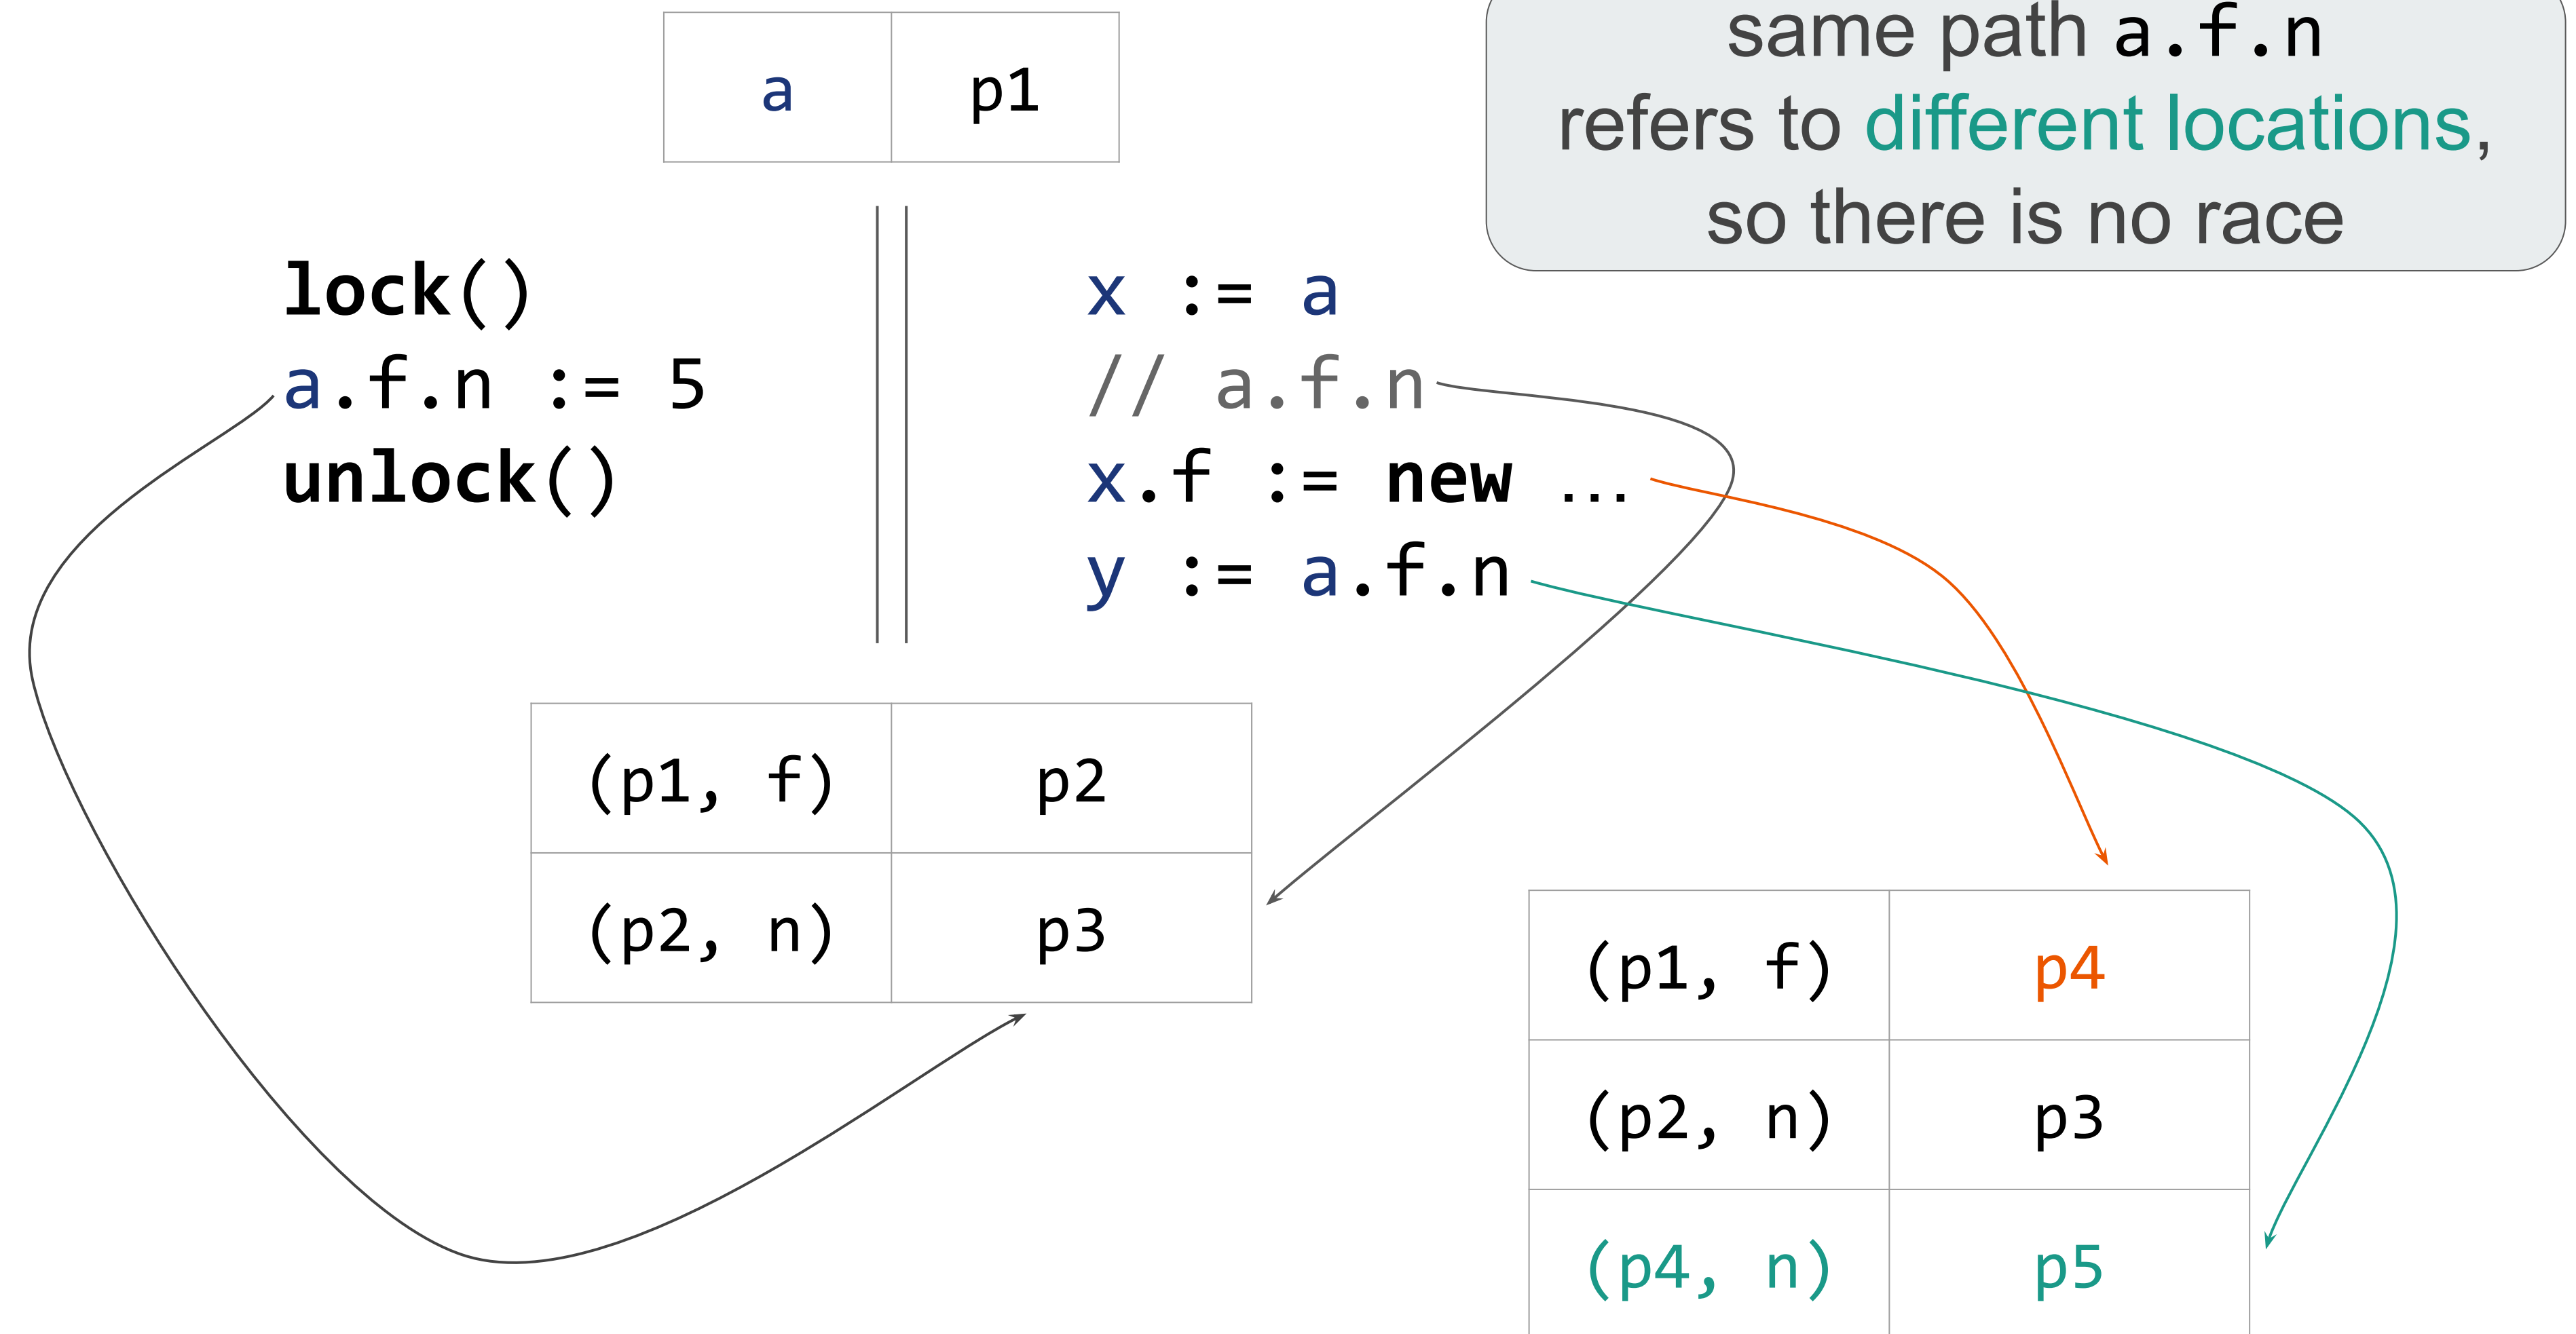
(1) abstract the beginning of the trace

abstract set of traces using exec

$\alpha(T) = \bigsqcup_{\tau \in T} (\text{fst } (\text{exec } \hat{\tau} (\perp_{\mathcal{D}}, \epsilon)))$

discard substitution

Why reading is wobbly?



Abstract access captures concrete access

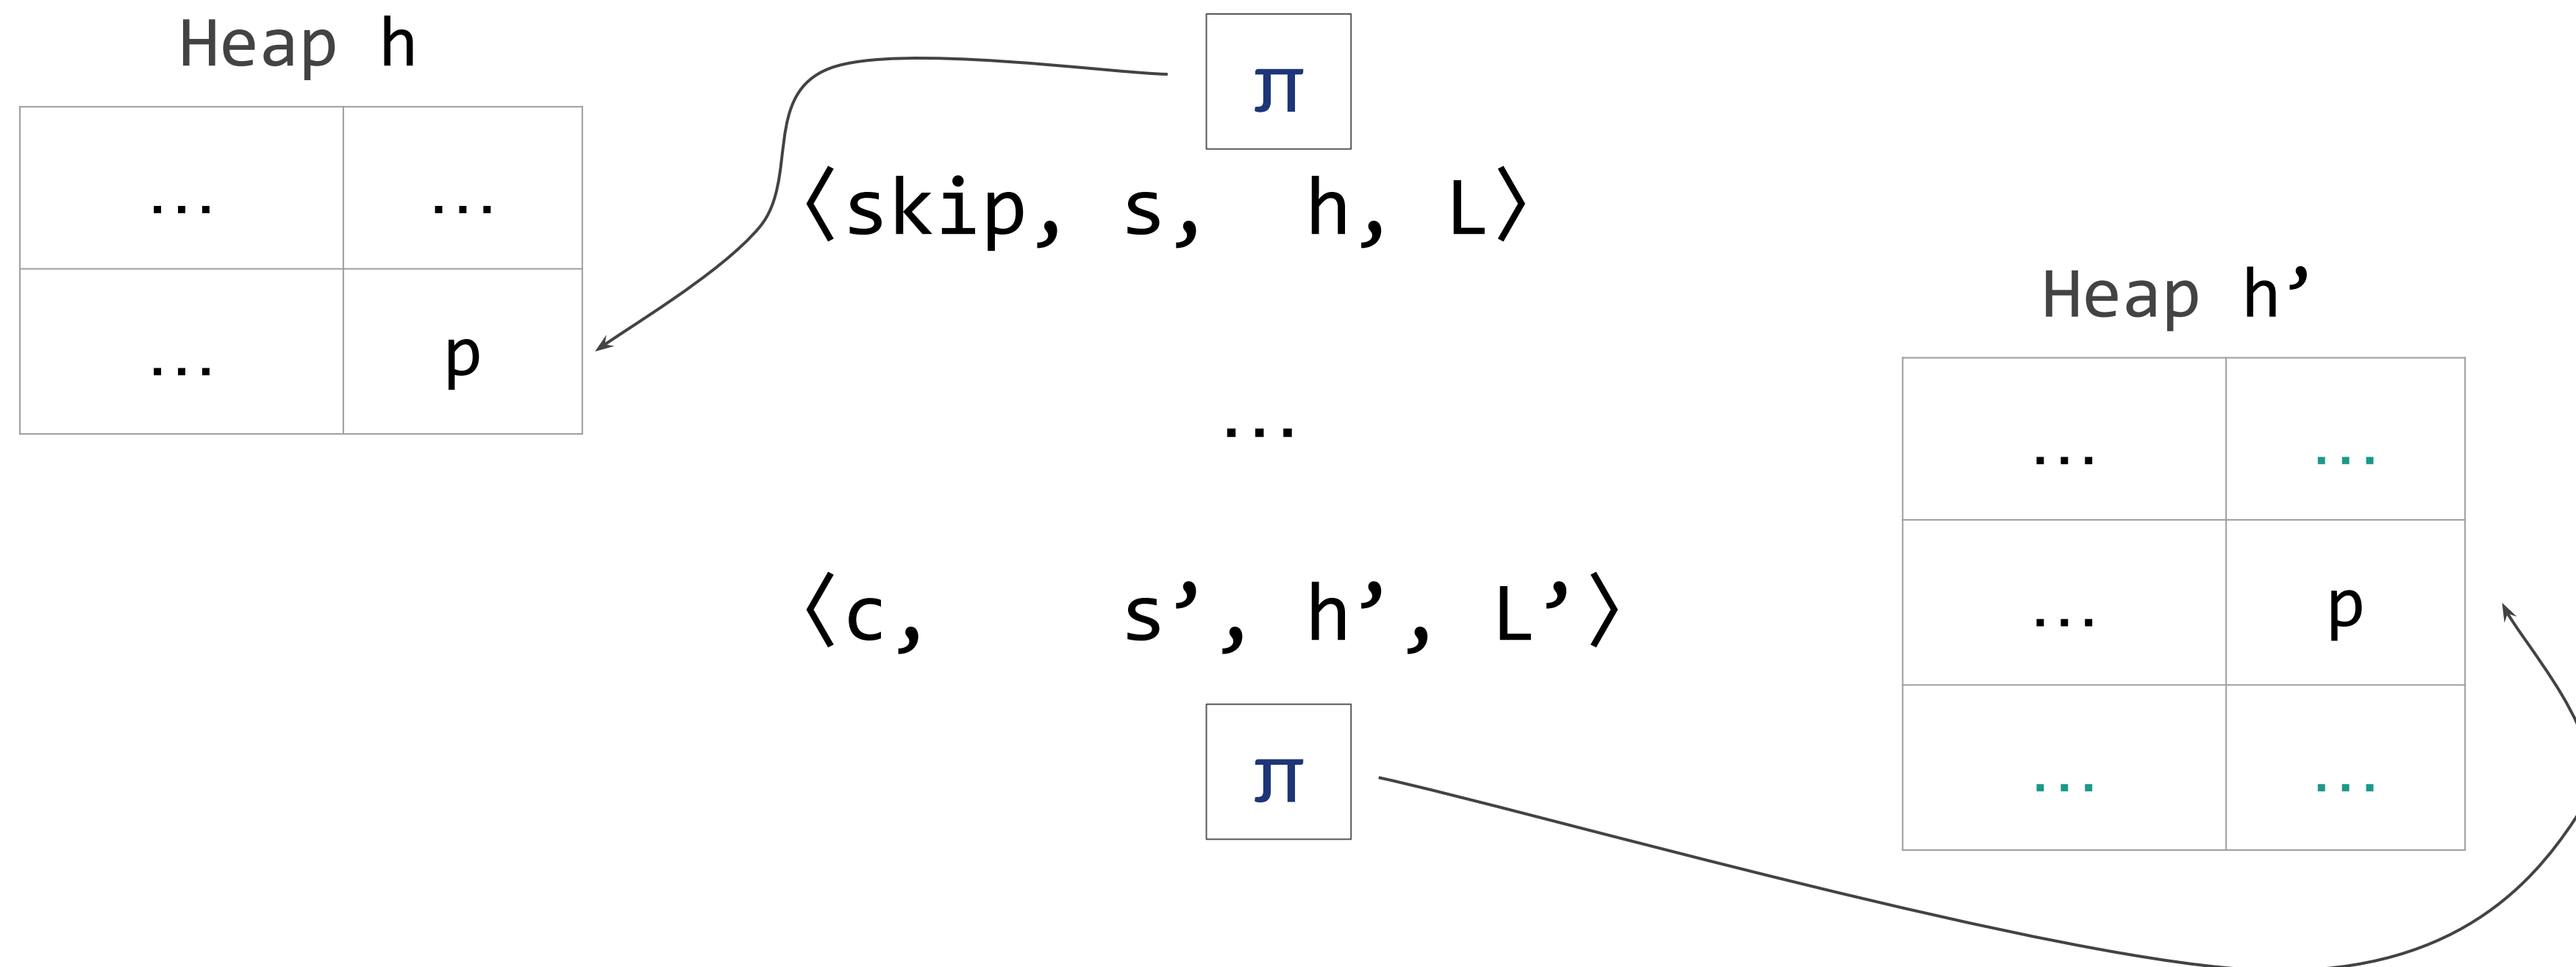
LEMMA 5.6 (PATH ACCESS EXISTENCE). *Let T be a set of traces, $\alpha(T) = \langle _, _, A \rangle$, and $q = \langle c, L \rangle$ (where $c = (x := \pi)$ or $c = (\pi := x)$) is a query about the access path π in the locking context L . If $q \in A$ then there exist a trace $\tau \in T$ and a non-empty, shortest prefix $\tau' \preceq \tau$ such that*

- *the last state of τ' is $\langle c', _, _, L \rangle$ and c', c are both stores or loads;*
- *$\text{exec } \widehat{\tau'} (\perp_{\mathcal{D}}, \epsilon) = (\langle _, _, A \rangle, \bar{\theta})$ where $\pi \in A$;*
- *a path π' such that $c' = (x := \pi')$ or $c' = (\pi' := x)$, and $\{\pi\} = \{\pi'\}^{\bar{\theta}}$.*

If a **path access** is recorded in the **abstract state**,
there is a **concrete trace** exhibiting the **access**

Stable path preserves memory location

If a good path is **not wobbly**,
it **preserves memory location** along a trace



Static Analysis

operates in
abstract domain

enjoys benefits
of the abstraction

- Does not need traces

$$\llbracket C; c \rrbracket^\# \langle W, L, A \rangle$$

- Compositional

$$\llbracket C; c \rrbracket^\# \langle W, L, A \rangle = \llbracket c \rrbracket^\# (\llbracket C \rrbracket^\# \langle W, L, A \rangle)$$

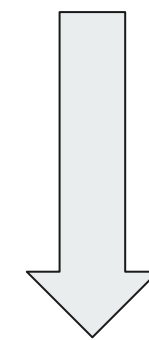
- **Complete** wrt. abstraction

$$\llbracket C \rrbracket^\# \perp_{\mathcal{D}} = \alpha (\llbracket C \rrbracket \langle S, h, 0 \rangle)$$

Galois connection

True Positives Theorem

$$\begin{array}{ccc} \llbracket C_1 \rrbracket^\# \perp_{\mathcal{D}} & \begin{array}{c} \text{access the same path } \pi \\ \text{in } W \parallel W \text{ or } W \parallel R \text{ or } R \parallel W \\ \text{and } \pi \text{ is not wobbly} \end{array} & \llbracket C_2 \rrbracket^\# \perp_{\mathcal{D}} \end{array}$$



$$\exists \varsigma^\parallel. \exists \text{racy } \tau^\parallel \in \llbracket C_1 \parallel C_2 \rrbracket^\# \varsigma^\parallel$$

π refers to the same location in C_1 & C_2 in the initial state, and it still refers to the same location when concurrently accessed

Evaluation

What is the price to pay for
having the TP Theorem?

(Reporting *no bugs whatsoever* is TP-Sound)

RacerD vs RacerDX

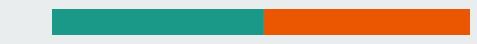
Target	LOC	D CPU	DX CPU	CPU $\pm\%$	D Reps	DX Reps	Reps $\pm\%$
avro	76k	103	102	0.4%	143	92	36%
Chronicle-Map	45k	196	196	0.1%	2	2	0%
jvm-tools	33k	106	109	-3.6%	30	26	13%
RxJava	273k	76	69	9.2%	166	134	19%
sunflow	25k	44	44	-1.4%	97	42	57%
xalan-j	175k	144	137	5.0%	326	295	10%

RacerD vs RacerDX

Target	LOC	D CPU	DX CPU	CPU $\pm\%$	D Reps	DX Reps	Reps $\pm\%$
avro	76k	103	102	0.4%	143	92	36%
Chronicle-Map	45k	196	196	0.1%	2	2	0%
jvm-tools	33k	106	109	-3.6%	30	26	13%
RxJava	273k	76	69	9.2%	166	134	19%
sunflow	25k	44	44	-1.4%	97	42	57%
xalan-j	175k	144	137	5.0%	326	295	10%

RacerD vs RacerDX

Target	LOC	D CPU	DX CPU	CPU $\pm\%$	D Reps	DX Reps	Reps $\pm\%$
avro	76k	103	102	0.4%	143	92	36%
Chronicle-Map	45k	196	196	0.1%	2	2	0%
jvm-tools	33k	106	109	-3.6%	30	26	13%
RxJava	273k	76	69	9.2%	166	134	19%
sunflow	25k	44	44	-1.4%	97	42	57%
xalan-j	175k	144	137	5.0%	326	295	10%



The Artifact

At Glance

Contents of Artifact



- Data: 6 packages source code
- Facebook's Infer package (OCaml code, Git repo): holds RacerD
- RacerDX Patch
- Set of Bash scripts to:
 - clean up
 - run vanilla Infer
 - patch and ran patched Infer,
 - collect stats.
- README: dependencies, entry points to run scripts, etc.



6 packages (incl. 2 invalid) by Build Technology

- Ant: 3 pkgs (avrora, sunflow, xalan-j)
- Gradle: 1 pkg (RxJava)
- Maven: 2 pkgs (Chronicle-Map, jvm-tools) — the invalid ones



Reproduction & reanalysis

War Stories

Repetition



- First try — failed with too new Java (noted in README):
 - error: as of release 9, '_' is a keyword, and may not be used as an identifier
- Second try — failed: unrecognized parameter to cloc (not noted in the README).
- Third try — partial success: numbers for RacerD are slightly off — ???
 - Reason: Missing native dependency
- Finally, numbers for 4 packages did reproduce.



What's Wrong with Maven?



Authors' words (README)

*Since submission of the paper for review, the sources of two of the projects (**Chronicle-Map** and **jvm-tools**) we used for evaluation became uncompileable (due to how maven works -- it **always** **downloads dependencies from the internet**, and it seems the newer versions are breaking the build of the version we originally tested).*

Dependencies? This should have to do with the **build scripts**!

pom.xml of Chronicle-map

```
28 <artifactId>chronicle-map</artifactId>
29 <version>3.16.0</version>
30 <name>OpenHFT/Chronicle-Map</name>
31 <description>Chronicle-Map</description>
32 <packaging>bundle</packaging>
33
34 <dependencyManagement>
35   <dependencies>
36
37     <dependency>
38       <groupId>net.openhft</groupId>
39       <artifactId>third-party-bom</artifactId>
40       <type>pom</type>
41       <version>3.6.5</version>
42       <scope>import</scope>
43     </dependency>
44
45     <dependency>
46       <groupId>net.openhft</groupId>
47       <artifactId>chronicle-bom</artifactId>
48       <version>1.16.134</version>
49       <type>pom</type>
50       <scope>import</scope>
```

Release

```
28 <artifactId>chronicle-map</artifactId>
29 <version>3.16.0-SNAPSHOT</version>
30 <name>OpenHFT/Chronicle-Map</name>
31 <description>Chronicle-Map</description>
32 <packaging>bundle</packaging>
33
34 <dependencyManagement>
35   <dependencies>
36
37     <dependency>
38       <groupId>net.openhft</groupId>
39       <artifactId>third-party-bom</artifactId>
40       <type>pom</type>
41       <version>3.6.2</version>
42       <scope>import</scope>
43     </dependency>
44
45     <dependency>
46       <groupId>net.openhft</groupId>
47       <artifactId>chronicle-bom</artifactId>
48       <version>LATEST</version>
49       <type>pom</type>
50       <scope>import</scope>
```

Artifact

Remember: “always downloads dependencies from the internet”...

pom.xml of jvm-tools (I)

```
1 <project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
2
3 <!--
4
5 Copyright 2012 Alexey Ragozin
6
7 Licensed under the Apache License, Version 2.0 (the "License");
8 you may not use this file except in compliance with the License.
9 You may obtain a copy of the License at
10
11 http://www.apache.org/licenses/LICENSE-2.0
12
13 Unless required by applicable law or agreed to in writing, software
14 distributed under the License is distributed on an "AS IS" BASIS,
15 WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
16 See the License for the specific language governing permissions and
17 limitations under the License.
18 -->
19
20
21 <modelVersion>4.0.0</modelVersion>
22
23 <parent>
24 <groupId>org.gridkit.lab</groupId>
25 <artifactId>grid-lab-pom</artifactId>
26 <version>2</version>
27 </parent>
28
29 <groupId>org.gridkit.jvmtool</groupId>
30 <artifactId>jvmtool-umbrella-pom</artifactId>
31 <version>0.11</version>
32 <name>${project.groupId}:${project.artifactId}</name>
```

Release

```
1 <?xml version="1.0" encoding="UTF-8"?>
2 <project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
3
4
5
6 <modelVersion>4.0.0</modelVersion>
7
8 <parent>
9 <groupId>org.gridkit.lab</groupId>
10 <artifactId>grid-lab-pom</artifactId>
11 <version>2</version>
12 </parent>
13
14 <groupId>org.gridkit.jvmtool</groupId>
15 <artifactId>jvmtool-umbrella-pom</artifactId>
16 <version>0.11-SNAPSHOT</version>
17 <name>${project.groupId}:${project.artifactId}</name>
```

Artifact

pom.xml of jvm-tools (II)



223 </project>

```
208 <profiles><profile>
209   <id>infer-capture</id>
210   <build>
211     <plugins>
212       <plugin>
213         <groupId>org.apache.maven.plugins</groupId>
214         <artifactId>maven-compiler-plugin</artifactId>
215         <configuration>
216           <compilerId>javac</compilerId>
217           <forceJavacCompilerUse>true</forceJavacCompilerUse>
218           <fork>true</fork>
219           <executable>/data/users/nikosgorogiannis/infer/infer/bin/infer</exe
220         </configuration>
221       </plugin>
222     </plugins>
223   </build>
224 </profile></profiles></project>
```



Infer's bug in management of pom.xml

- Error message like the one in the artifact (when enabling stderr):

```
Error while running epilogue restoring Maven's pom.xml to its original state:
(Unix.Unix_error "No such file or directory"
  rename "((src /data/videoRecorder/videoRecorder-rpm/pom.xml.infer-orig)
    (dst /data/videoRecorder/videoRecorder-rpm/pom.xml)))".
```

- There's a fix also!

The Fix

facebook / infer

Watch 562 Star 10.3k Fork 1.4k

Code Issues 213 Pull requests 6 Actions Projects 0 Wiki Security Insights

✓ [epilogues] do not rely on `at_exit`

Browse files

Summary:

For some unexplained reason, some of the functions registered in the Epilogues would sometimes be executed several times. I could not figure out why.

This diff fixes that, but also has more explainable benefits:

- Do not run epilogues registered in the parent in the children. Previously it would do so, but probably only if the children registered some epilogue given that `at_exit` must be called again once on the child (but the value of the ref in `Pervasives` would not have been reset).
- Unified behaviour for early and late epilogues given that we now handle both of these directly

We already have all the control needed to run epilogues when needed: we know when infer exits, and we know when children processes exit.

Reviewed By: mbouaziz

Differential Revision: D9752046

fbshipit-source-id: 13af40081

master v0.17.0 v0.16.0

jvillard

 authored and facebook-github-bot committed on Sep 12, 2018

1 parent 4ddbc71 commit 817f83972cacba03bd9613979fd55154a6b95e20



Happy End with Repetition

After

- Applying the Infer fix (*kudos to the authors for preserving the Git repo*)
- Checking out released versions of Maven-based packages

We were able to get the numbers from the paper.



Our Experiments (mostly support the claims)

- Full aws-sdk-java died with disk overflow (hundreds of gigabytes of reports)
 - Just one module (aws-java-sdk-s3):
 - test-aws-sdk-java, 3'847'035, 666, 639, 64, 48
- spring-kafka — success, equal results:
 - test-spring-kafka, 30'461, 31, 31, 16, 16
- azkaban — success, equal, zeroes:
 - test-azkaban, 76'156, 0, 0, 0, 0


```
{
  "bug_class": "PROVER",
  "kind": "ERROR",
  "bug_type": "THREAD_SAFETY_VIOLATION",
  "qualifier": "Unprotected write. Non-private method `void EmbeddedKafkaBroker.afterPropertiesSet()` writes to field `this.org.springframework.kafka.test.EmbeddedKafkaBroker.zkConnect` outside of synchronization
  .\n Reporting because another access to the same memory occurs on a background thread, although this access may not.",
  "severity": "HIGH",
  "visibility": "user",
  "line": 267,
  "column": -1,
  "procedure": "void EmbeddedKafkaBroker.afterPropertiesSet()",
  "procedure_id": "org.springframework.kafka.test.EmbeddedKafkaBroker.afterPropertiesSet():void.9fb246b16e74937d3be8f9fd6f1d35ed",
  "procedure_start_line": 0,
  "file": "spring-kafka-test/src/main/java/org/springframework/kafka/test/EmbeddedKafkaBroker.java",
  "bug_trace": [
    {
      "level": 0,
      "filename": "spring-kafka-test/src/main/java/org/springframework/kafka/test/EmbeddedKafkaBroker.java",
      "line_number": 267,
      "column_number": -1,
      "description": "<Write on unknown thread>"
    },
    {
      "level": 0,
      "filename": "spring-kafka-test/src/main/java/org/springframework/kafka/test/EmbeddedKafkaBroker.java",
      "line_number": 267,
      "column_number": -1,
      "description": "access to `this.org.springframework.kafka.test.EmbeddedKafkaBroker.zkConnect`"
    },
    {
      "level": 0,
      "filename": "spring-kafka-test/src/main/java/org/springframework/kafka/test/EmbeddedKafkaBroker.java",
      "line_number": 415,
      "column_number": -1,
      "description": "<Write on background thread>"
    },
    {
      "level": 0,
      "filename": "spring-kafka-test/src/main/java/org/springframework/kafka/test/EmbeddedKafkaBroker.java",
      "line_number": 415,
      "column_number": -1,
      "description": "access to `this.org.springframework.kafka.test.EmbeddedKafkaBroker.zkConnect`"
    }
  ],
  "key": "EmbeddedKafkaBroker.java|afterPropertiesSet|THREAD_SAFETY_VIOLATION",
  "node_key": "9c5d6d9028928346cc4fb44cced5deal",
  "hash": "69130a79005ccdc6e56b21690faa2fla",
  "bug_type_hum": "Thread Safety Violation",
  "censored_reason": "",
  "access": "hJWmvGAAAdwAAAAzAAAA5QAAALSwkNAyYWZ0ZXJQcm9wZXJ0aWVzU2V0QJ0gCTJvcmcuc3ByaW5nZnJhbWV3b3JrLmthZmthLnRlc3QuRW1iZWRLZWRLYWZrYUJyb2t1ckCQoEAKdm9pZECgoJKgIECwAQELAP
+SCVdzchJpbmcta2Fma2EtdGVzdC9zcmMvbWVpbi9qYXZhL29yZy9zcHJpbmdmcmFtZXdvcmstva2Fma2EvdGVzdC9FbWJlZGRlZEthZmthQnJva2VyLmphdmGRoKCRsAI2n0wXoCR0aGlzQJCQ0DJhZnRlc1Byb3BlcnRpZXNTZXRAK6AJMm9yZy5zcHJpbmdmcmFtZXdvcmstva2Fma2EudGVzdC5FbWJlZGRlZEthZmthQnJva2VyLnprQ29ubmVjdECgBCCgsAEBnwD
/kglXc3ByaW5nLWthZmthLXRlc3Qvc3JjL2lhaW4vamF2YS9vcmcvc3ByaW5nZnJhbWV3b3JrL2thZmthL3Rlc3QvRW1iZWRLZWRLYWZrYUJyb2t1ci5qYXZhQA=="
},
```

Race Report Example

RacerDX

```
1 {
2   "bug_class": "PROVER",
3   "kind": "ERROR",
4   "bug_type": "THREAD_SAFETY_VIOLATION",
5   "qualifier": "Read/Write race. Non-private method `JPanel GraphNumbers.chalk
boardAndBar()` indirectly reads with synchronization from `this.avrora.gui.Graph
Numbers.privateNumbers.cck.stat.Sequence.total`. Potentially races with unsynchr
onized write in method `GraphNumbers.internalUpdate()`. \n Reporting because anot
her access to the same memory occurs on a background thread, although this acces
s may not.",
6   "severity": "HIGH",
7   "visibility": "user",
8   "line": 123,
9   "column": -1,
10  "procedure": "JPanel GraphNumbers.chalkboardAndBar()",
11  "procedure_id": "avrora.gui.GraphNumbers.chalkboardAndBar():javax.swing.JPan
el.ecda2ba0ceefadb1302bcf0ca8e35bca",
12  "procedure_start_line": 0,
13  "file": "src/avrora/gui/GraphNumbers.java",
14  "bug_trace": [
15    {
16      "level": 0,
17      "filename": "src/avrora/gui/GraphNumbers.java",
18      "line_number": 123,
19      "column_number": -1,
20      "description": "<Read trace>"
21    },
22    {
23      "level": 0,
24      "filename": "src/avrora/gui/GraphNumbers.java",
25      "line_number": 123,
26      "column_number": -1,
27      "description": "call to void GraphNumbers.updateHorzBar()"
28    },
29    {
30      "level": 1,
31      "filename": "src/avrora/gui/GraphNumbers.java",
32      "line_number": 151,
33      "column_number": -1,
34      "description": "call to int Sequence.size()"
35    },
36    {
37      "level": 2,
38      "filename": "src/cck/stat/Sequence.java",
39      "line_number": 157,
40      "column_number": -1,
41      "description": "access to `this.cck.stat.Sequence.total`"
42    },
43  ]
44 }
```

RacerD

```
1 {
2   "bug_class": "PROVER",
3   "kind": "ERROR",
4   "bug_type": "THREAD_SAFETY_VIOLATION",
5   "qualifier": "Read/Write race. Non-private method `JPanel GraphNumbers.chalk
boardAndBar()` indirectly reads with synchronization from `this.avrora.gui.Graph
Numbers.privateNumbers.cck.stat.Sequence.total`. Potentially races with unsynchr
onized write in method `GraphNumbers.internalUpdate()`. \n Reporting because anot
her access to the same memory occurs on a background thread, although this acces
s may not.",
6   "severity": "HIGH",
7   "visibility": "user",
8   "line": 123,
9   "column": -1,
10  "procedure": "JPanel GraphNumbers.chalkboardAndBar()",
11  "procedure_id": "avrora.gui.GraphNumbers.chalkboardAndBar():javax.swing.JPan
el.ecda2ba0ceefadb1302bcf0ca8e35bca",
12  "procedure_start_line": 0,
13  "file": "src/avrora/gui/GraphNumbers.java",
14  "bug_trace": [
15    {
16      "level": 0,
17      "filename": "src/avrora/gui/GraphNumbers.java",
18      "line_number": 123,
19      "column_number": -1,
20      "description": "<Read trace>"
21    },
22    {
23      "level": 0,
24      "filename": "src/avrora/gui/GraphNumbers.java",
25      "line_number": 123,
26      "column_number": -1,
27      "description": "call to void GraphNumbers.updateHorzBar()"
28    },
29    {
30      "level": 1,
31      "filename": "src/avrora/gui/GraphNumbers.java",
32      "line_number": 151,
33      "column_number": -1,
34      "description": "call to int Sequence.size()"
35    },
36    {
37      "level": 2,
38      "filename": "src/cck/stat/Sequence.java",
39      "line_number": 157,
40      "column_number": -1,
41      "description": "access to `this.cck.stat.Sequence.total`"
42    },
43  ]
44 }
```

Report Subsets



Takeaways: How NOT To Make an Artifact

- No Environment Management (e.g. a VM, Docker, Nix, etc.):
 - a. a bunch of source codes (sometimes non-released versions; not tracked by a VCS, e.g. Git)
 - b. (lose) description of dependencies (some dependencies didn't have corresponding versions, e.g. cloc)
 - c. No way to account for transitive deps of tools, esp. native deps (e.g. sqlite3-dev)
- Clearing the \$PATH -- poor man's env management
- Piping stdout and stderr (!!!) to /dev/null