

Announcement:
The lesson scheduled for
Friday, the 24th, will not
be held



Course Focus and Next Steps

In this course, we have explored how to model systems and then verify them using model checking techniques.

This represents the main line of the course.

Now, we will take a side path and examine some program analysis techniques.

We will focus on dynamic analysis techniques.

What is program analysis?

Testing and Analysis

What you know
from previous
courses

- Manual testing or semi-automated testing
- JUnit, Pytest, Selenium, etc.

Manual
“analysis” of
programs:

- Code inspection, debugging, etc.

```
import org.junit.jupiter.api.*;
import static org.junit.jupiter.api.Assertions.*;

class GeneralTestSuite {

    private ExampleService service;

    @BeforeAll
    static void initAll() {
        System.out.println("Starting test suite...");
    }

    @BeforeEach
    void init() {
        service = new ExampleService();
    }

    @Test
    void testAddition() {
        assertEquals(4, service.add(2, 2), "Addition should return correct result");
    }
}
```

```
@Test
void testDivisionByZeroThrows() {
    assertThrows(ArithmeticException.class, () -> service.divide(10, 0));
}

@Test
void testStringNotNull() {
    String result = service.getMessage();
    assertNotNull(result, "Message should not be null");
    assertTrue(result.startsWith("Hello"), "Message should start with 'Hello'");
}

@AfterEach
void tearDown() {
    service = null;
}

@AfterAll
static void tearDownAll() {
    System.out.println("All tests completed.");
}
```

```
// Dummy implementation under test
class ExampleService {
    int add(int a, int b) { return a + b; }
    int divide(int a, int b) { return a / b; }
    String getMessage() { return "Hello, World!"; }
}
```

The Key Concepts



@BeforeAll / @AfterAll for global setup and cleanup



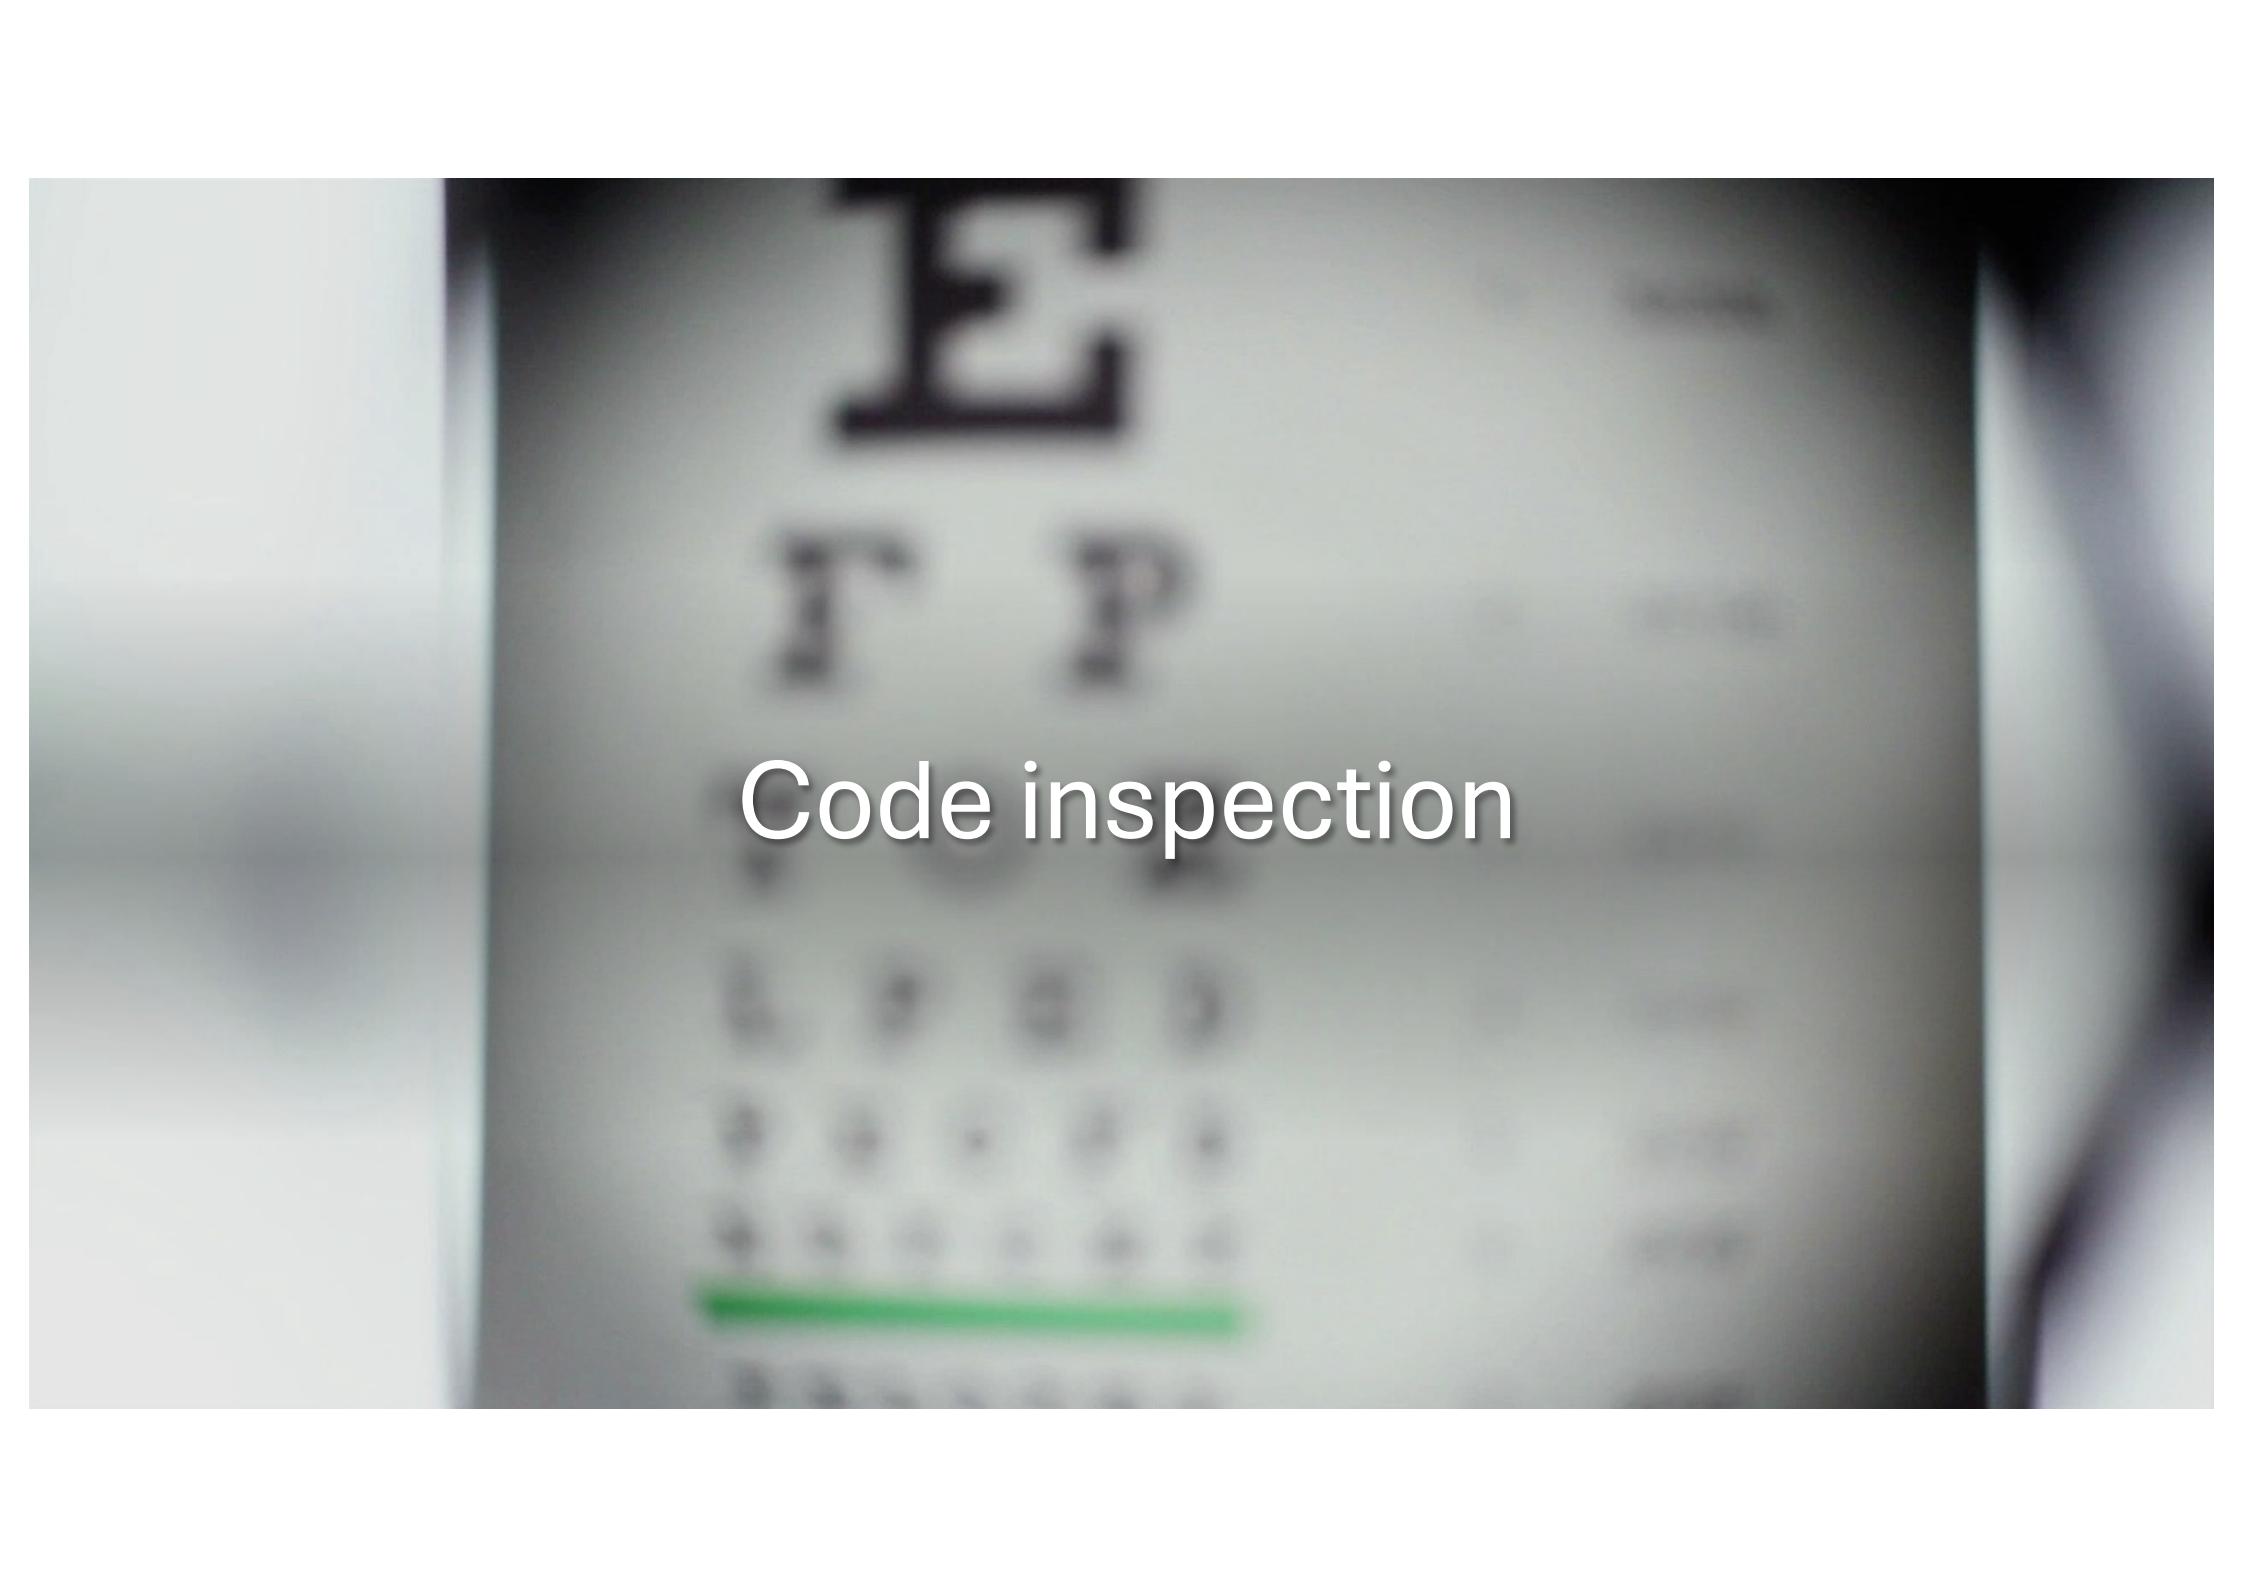
@BeforeEach / @AfterEach for per-test setup and teardown



Assertions: assertEquals, assertNotNull, assertThrows ...



Good naming and descriptive messages

A blurry background image of a computer screen showing code in a terminal window. The text 'E' and 'T P' are visible at the top, and a green horizontal bar is at the bottom.

Code inspection

Rust Code

```
use std::env;

fn read_port() -> u16 {
    let port_str = env::var("APP_PORT").unwrap();
    let port: i32 = port_str.parse().unwrap();
    port as u16
}
```

The function `read_port()` reads the environment variable `APP_PORT`, parses it as a number, and returns it as a port.

Rust Code

```
use std::env;

fn read_port() -> u16 {
    let port_str = env::var("APP_PORT").unwrap();
    let port: i32 = port_str.parse().unwrap();
    port as u16
}
```

The function `read_port()` reads the environment variable `APP_PORT`, parses it as a number, and returns it as a port.

`unwrap()` is a method used on `Option<T>` and `Result<T, E>` types. It retrieves the inner value but will `panic` if the value is `None` or `Err`.

```
use std::env;

fn read_port() -> u16 {
    let port_str = env::var("APP_PORT").unwrap(); // panics if missing
    let port: i32 = port_str.parse().unwrap(); // panics if invalid
    port as u16
}
```

```
use std::env, num::ParseIntError;

#[derive(Debug)]
enum PortError {
    MissingVar,
    Parse(ParseIntError),
}

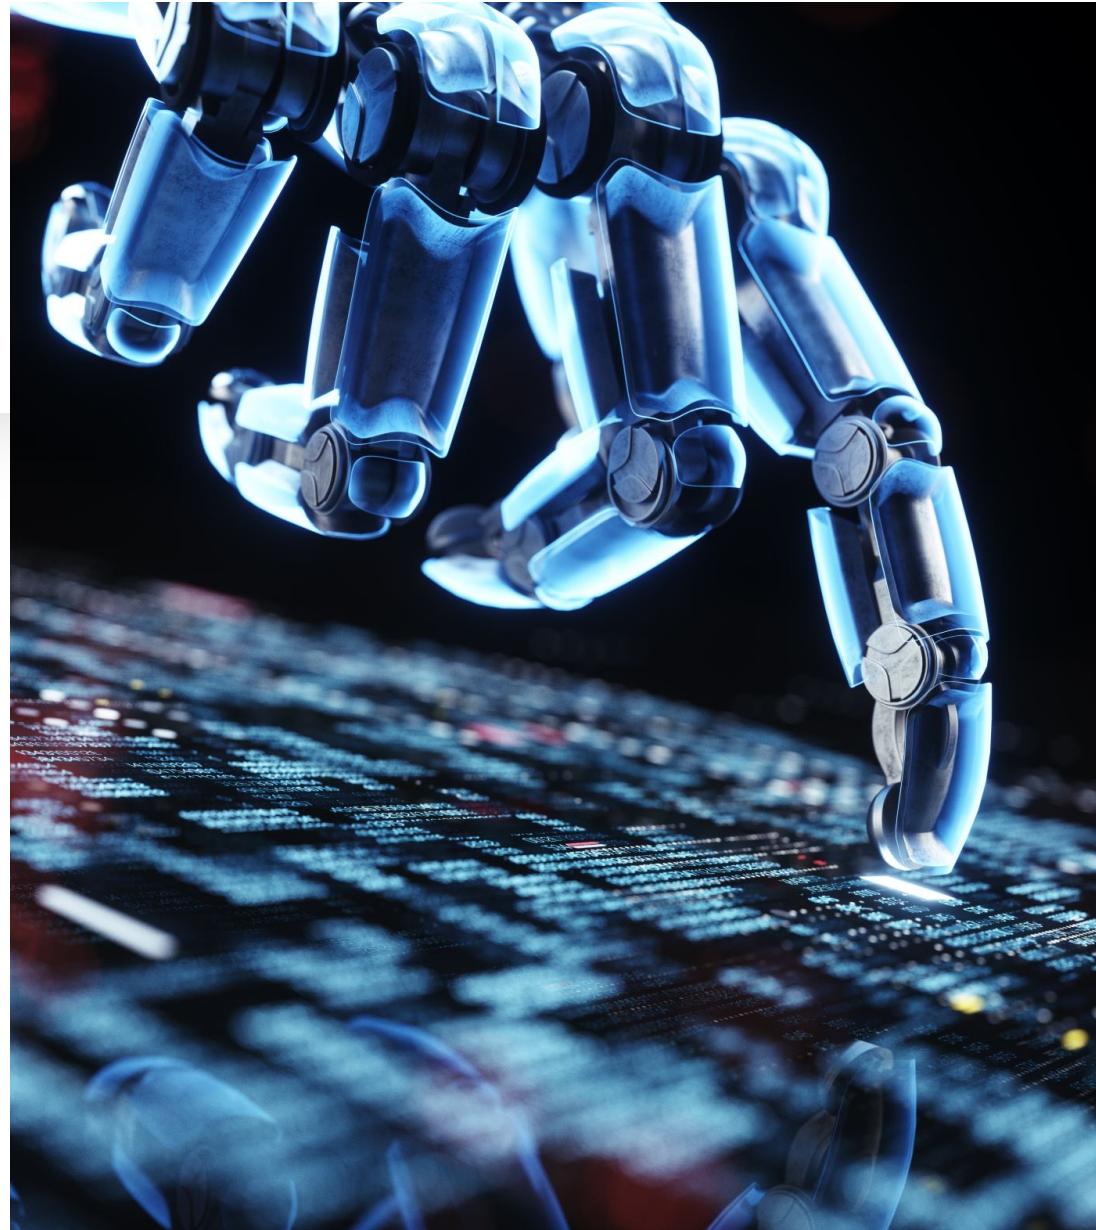
fn read_port() -> Result<u16, PortError> {
    let raw = env::var("APP_PORT").map_err(|_| PortError::MissingVar)?;
    let port = raw.parse::<u16>().map_err(PortError::Parse)?;
    Ok(port)
}
```



Focus of this side path of the course:
Automated testing and
program analysis

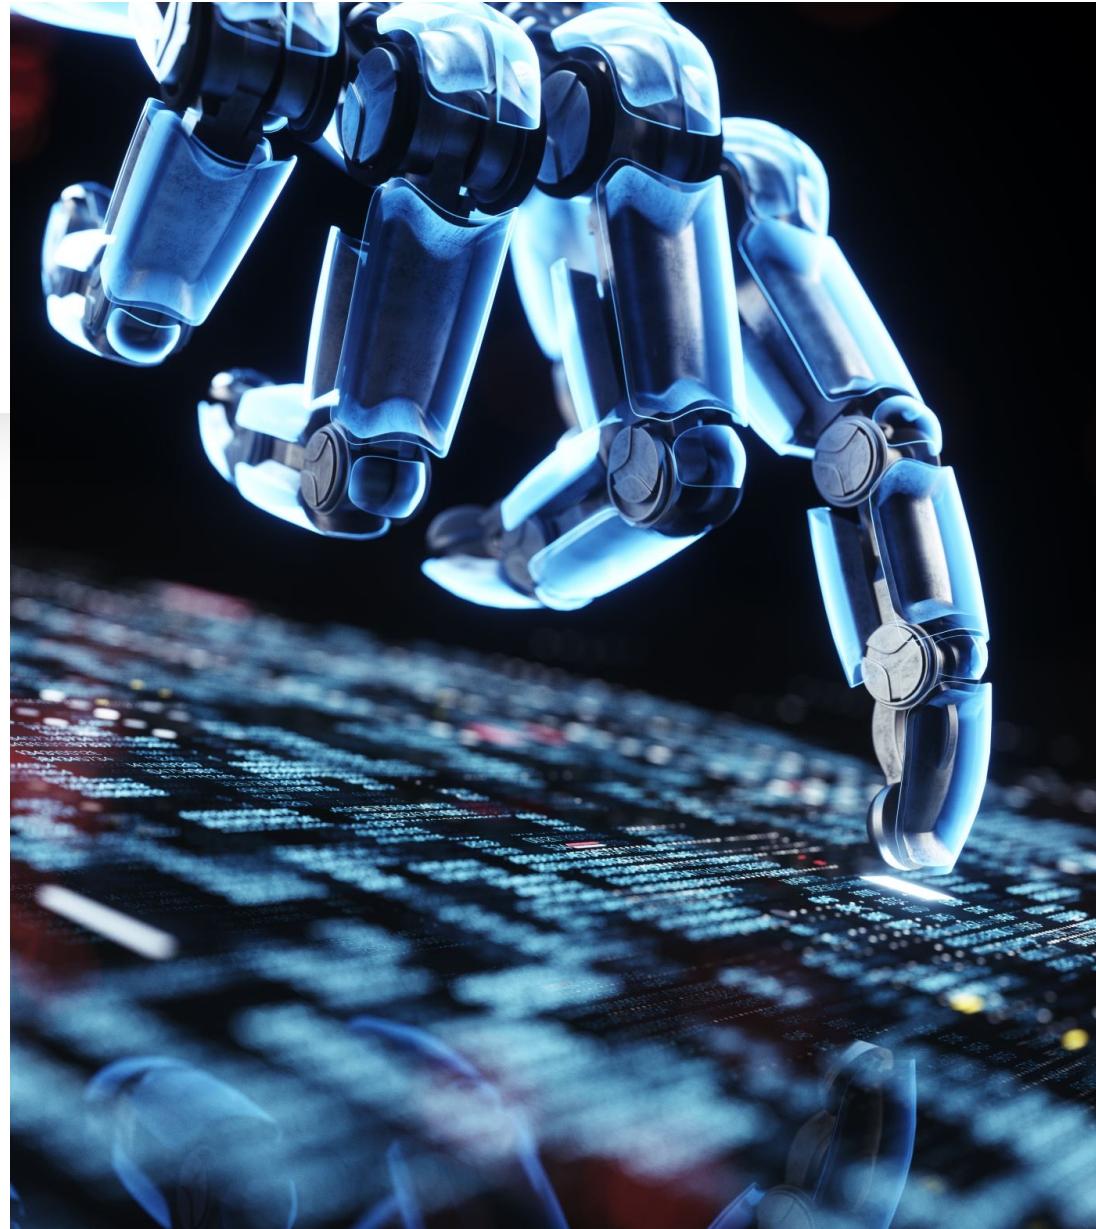
Why do we talk about program analysis and automated testing?

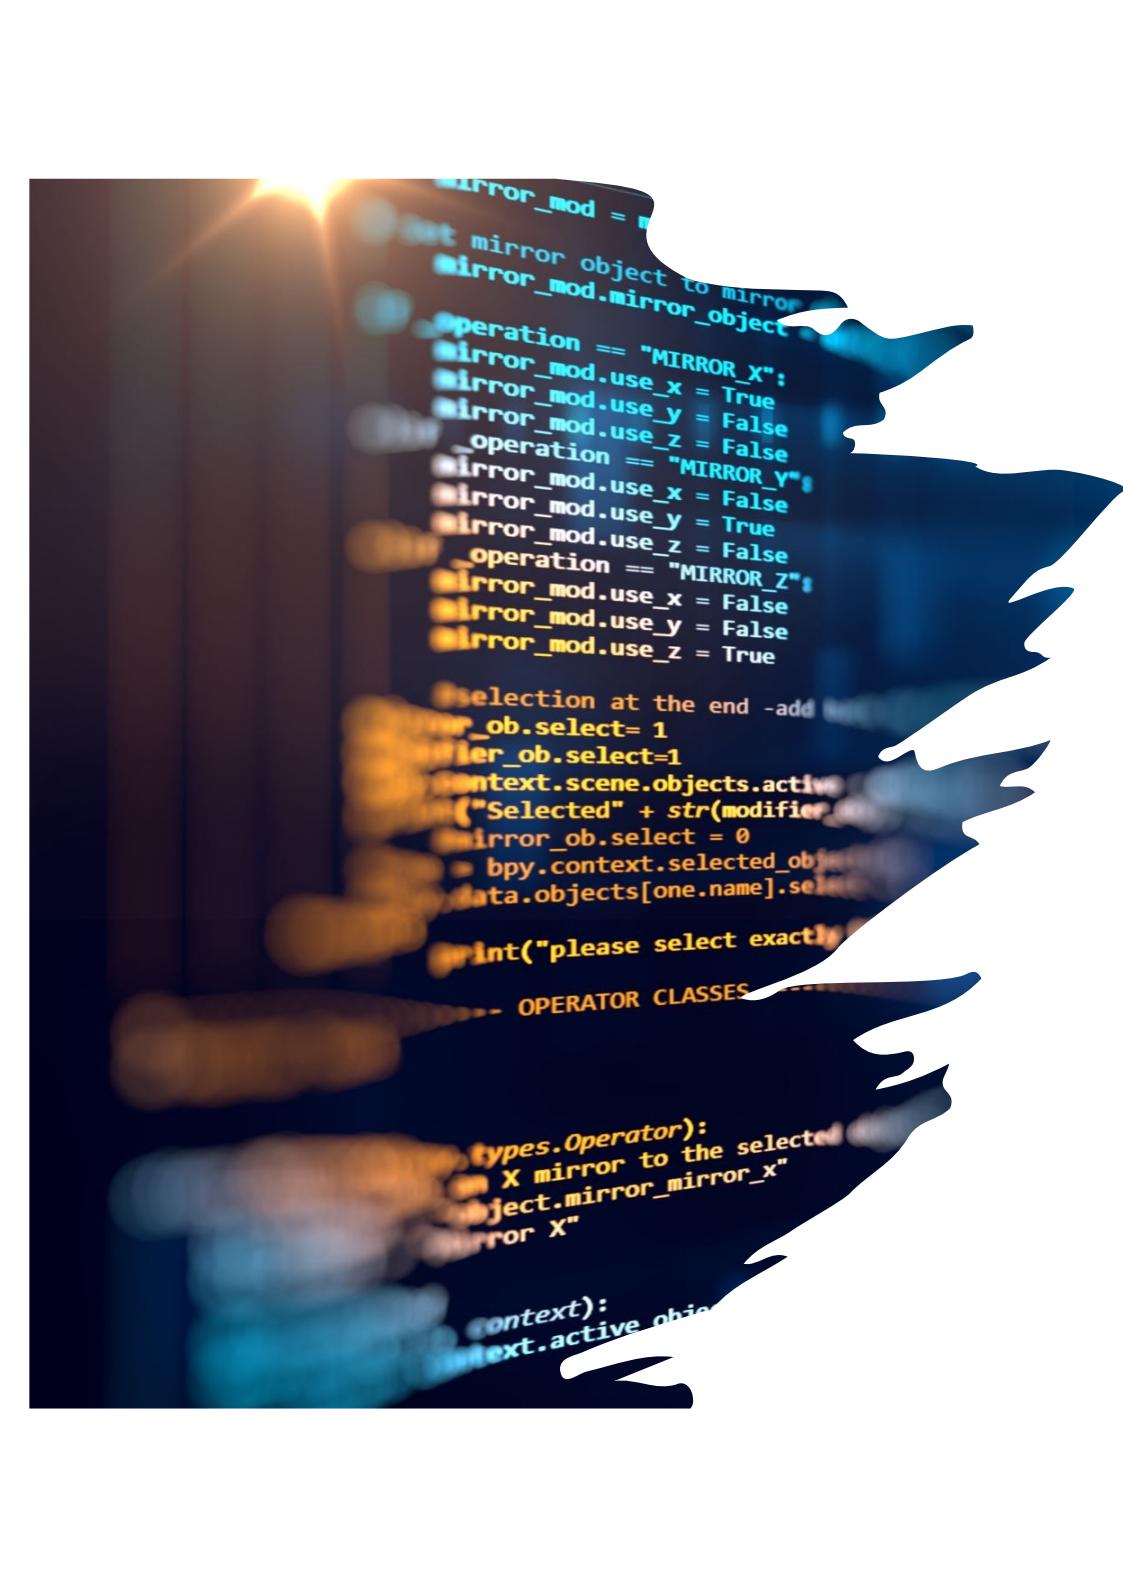
- All software has bugs
- Bugs are hard to find
- Bugs cause serious harm



Why do we talk about program analysis and automated testing?

- All software has bugs
- Bugs are hard to find
- Bugs cause serious harm
- ... Because no software is bug-free, an AI-based copilot can help us to spot and resolve bugs more efficiently.





My personal statement

- While AI-based copilots offer valuable support in software development, human understanding of the generated code remains essential, as software (**even AI-generated software**) may still contain bugs.

What is program analysis

Automated analysis of program behavior capable to

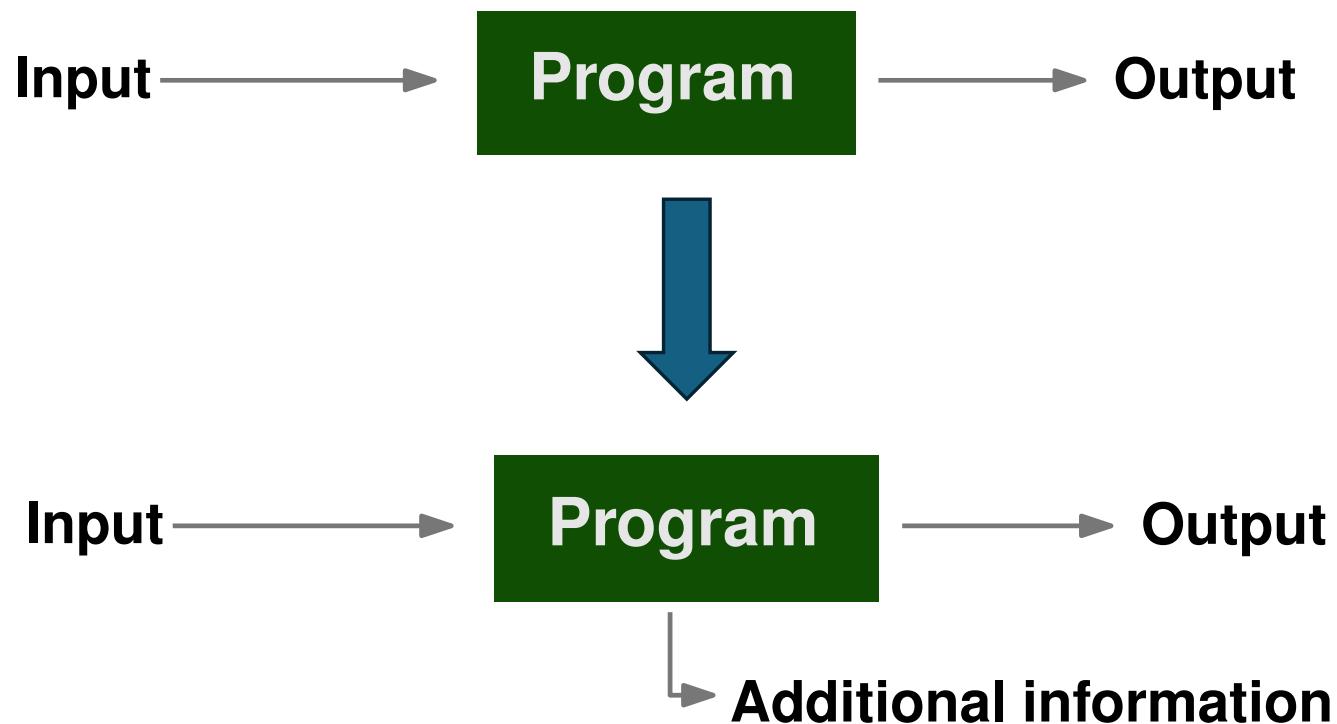
find programming errors

optimize performance

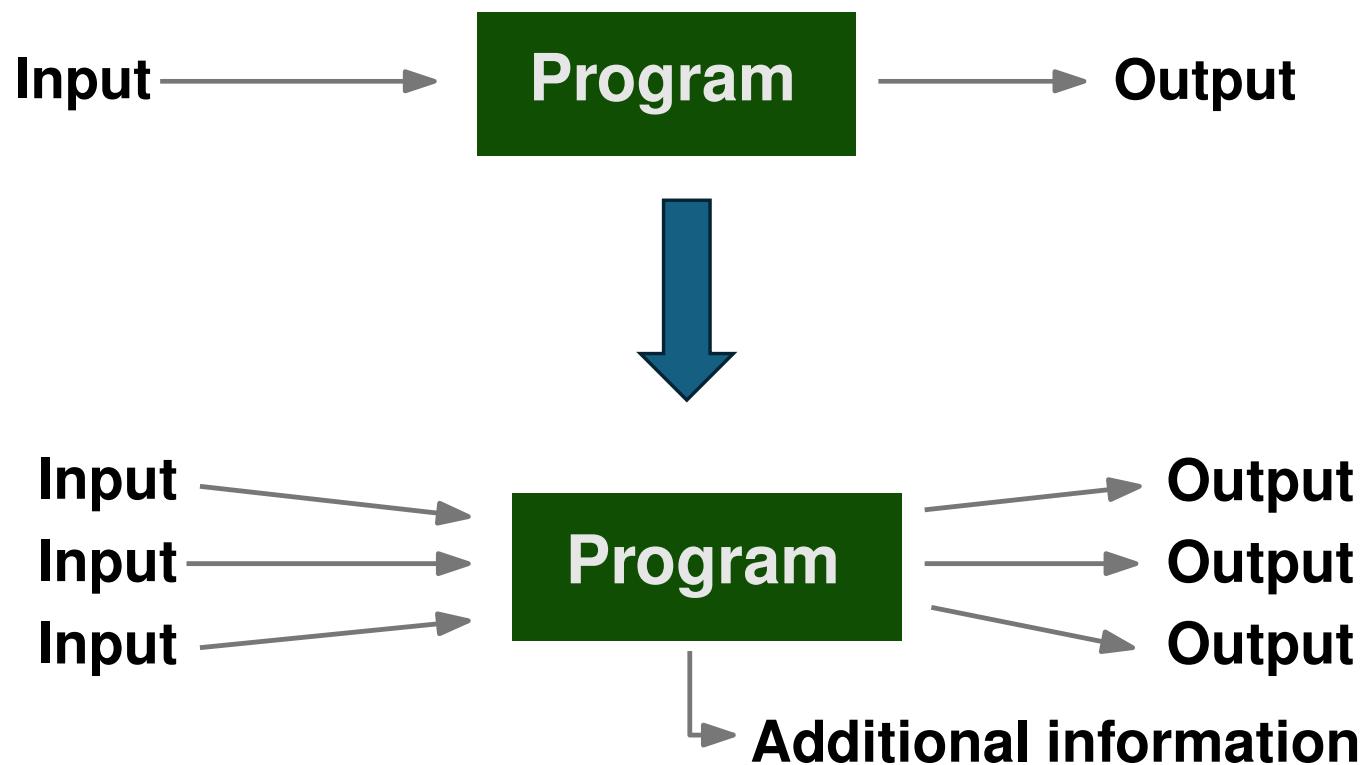
find security vulnerabilities

...

From Program Execution to Program Analysis



Program Analysis



Program analysis

Static

- Analyze source code, byte code, or binary
- Typically:
 - Consider all inputs
 - **Overapproximate possible behavior**

Dynamic

- Analyze program execution
- Typically:
 - Consider current input
 - **Underapproximate possible behavior**

Program analysis example

```
// Javascript
var r = Math.random();
var out = "yes";

if (r < 0.5)
    out = "no";

if (r === 1)
    out = "maybe";

console.log(out);
```

Program analysis example

```
// Javascript
var r = Math.random();
var out = "yes";

if (r < 0.5)
    out = "no";

if (r === 1)
    out = "maybe";

console.log(out);
```

Overapproximation:
"yes", "no", "maybe"

Consider all paths

Program analysis example

```
// Javascript
var r = Math.random();
var out = "yes";

if (r < 0.5)
    out = "no";

if (r === 1)
    out = "maybe";

console.log(out);
```

Underapproximation:
"yes"

Execute the program once

Program analysis example

```
// Javascript
var r = Math.random();
var out = "yes";

if (r < 0.5)
    out = "no";

if (r === 1)
    out = "maybe";

console.log(out);
```

Sound and complete analysis::
"yes", "no"

The analysis explores both feasible paths

Program analysis example

```
// Javascript
var r = Math.random();
var out = "yes";

if (r < 0.5)
    out = "no";

if (r === 1)
    out = "maybe";

console.log(out);
```

Math.random()

Returns a floating-point number r such that $0 \leq r < 1$
It never returns exactly 1.

Initialization

out starts as "yes".

First if condition ($r < 0.5$)

If the random number is less than 0.5, out becomes "no".
This happens roughly 50% of the time.

Second if condition ($r === 1$)

This would only run if r is exactly 1.
But since `Math.random()` never produces 1, this branch
is impossible

Output

So out is either "no" or "yes" depending on r .

Another example

```
//Jvascript
var r = Math.random();

var out = r*2

console.log(out);
```

Another example

```
// Javascript
var r = Math.random();

var out = r*2

console.log(out);
```

Method	Output	Discussion
Over-approximation		
Under-approximation		
Sound and Complete		

Another example

```
// Javascript
var r = Math.random();

var out = r*2

console.log(out);
```

Method	Output	Discussion
Over-approximation	Any value	All paths
Under-approximation		
Sound and Complete		

Another example

```
// Javascript
var r = Math.random();

var out = r*2

console.log(out);
```

Method	Result	Discussion
Over-approximation	Any value	All paths
Under-approximation	Some number in [0,2), e.g., 1.234	One execution
Sound and Complete		

Another example

```
// Javascript
var r = Math.random();

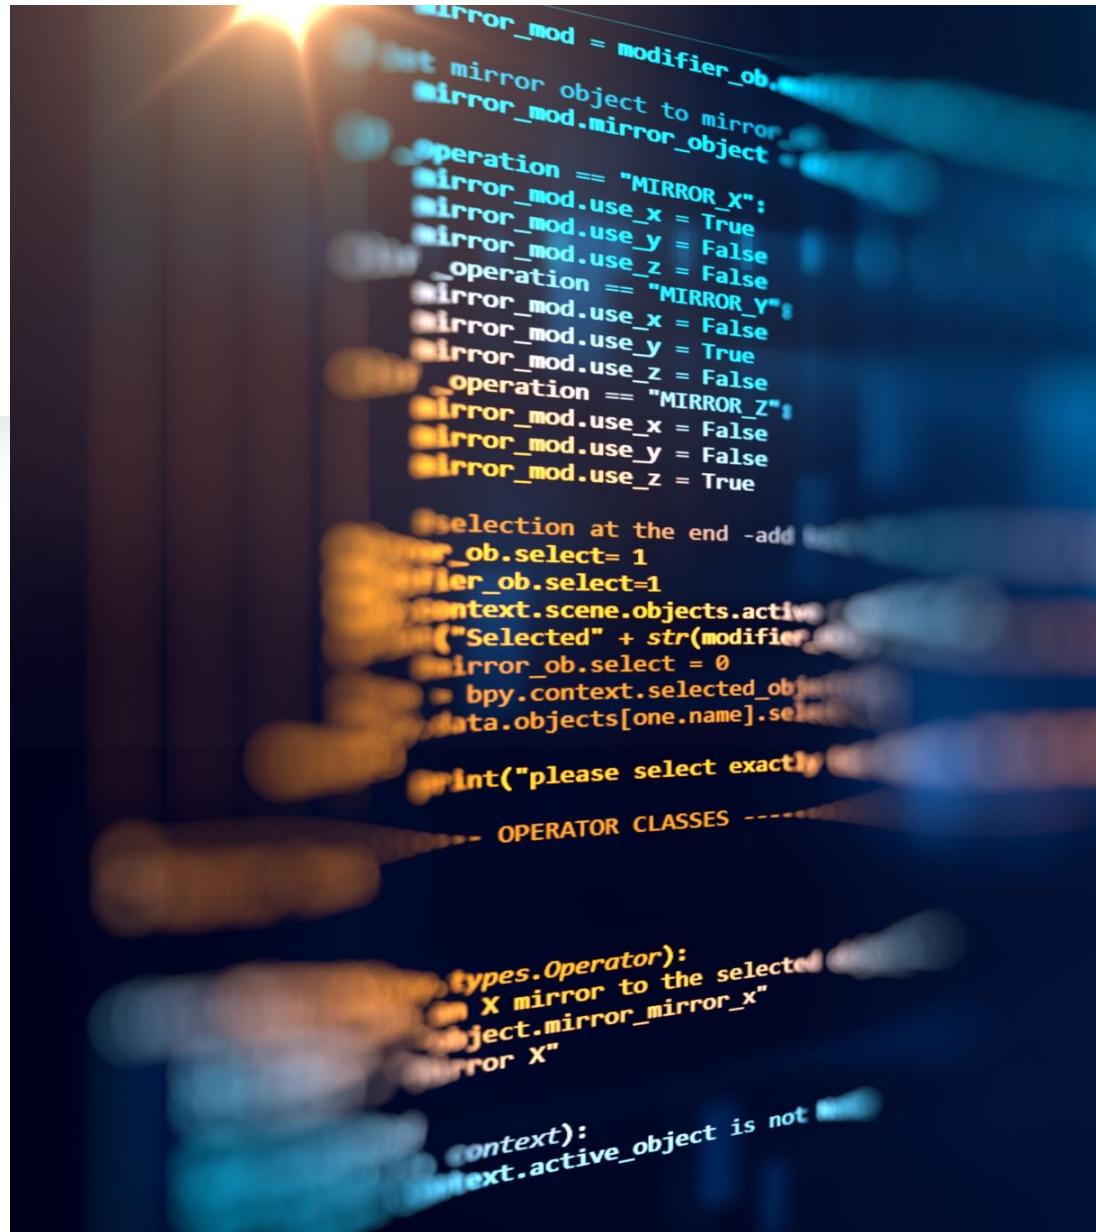
var out = r*2

console.log(out);
```

Method	Result	Discussion
Over-approximation	Any value	All paths
Under-approximation	Some number in [0,2), e.g., 1.234	One execution
Sound and Complete	??	Exploring all possible outputs: Practically impossible

How does program analysis help me?

- **Use program analysis tools**
 - Improve the quality of your code
- **Understand program analysis**
 - Better understanding of program code and program behaviours
- **Create your own tools**



Dynamic Analysis

Dynamic Analysis

Execute an **instrumented program** to gather information that can be analyzed to learn about a property of interest

Precise: All observed behavior actually happens

Incomplete: Very difficult to cover all possible behaviors

Dynamic Analysis Examples

Coverage: Track which lines or branches get executed

Call graph: Track which functions call which other functions

Slicing: Track dependencies to produce a reduced program

Some references

- Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation, Nethercote et al., PLDI 2007
- Jalangi: A Selective Record-Replay and Dynamic Analysis Framework for JavaScript, Sen et al., FSE 2013
- DynaPyt: A Dynamic Analysis Framework for Python, Eghbali et al., FSE 2022

Call graph analysis

```
function main() {  
    let n = readInput();  
  
    function a() {  
        b();  
    }  
  
    function b() {  
        if (n == 5) {  
            c();  
        }  
    }  
  
    function c() {  
        if (n == 5) {  
            c();  
            n--;  
        }  
    }  
  
    a();  
}
```

Call graph analysis

Nodes

main, readInput, a, b, c

Edges

main → readInput

main → a

a → b

b → c (*if n == 5*)

c → c (*if n == 5*). // **recursive call**

```
function main() {  
    let n = readInput();  
  
    function a() {  
        b();  
    }  
  
    function b() {  
        if (n == 5) {  
            c();  
        }  
    }  
  
    function c() {  
        if (n == 5) {  
            n--;  
            c();  
        }  
    }  
  
    a();  
}
```

The instrumented code

```
function main() {  
  const calls = new Set();  
  
  let n = readInput();  
  calls.add("main->readInput");  
  
  function a() {  
    calls.add("a->b");  
    b();  
  }  
  
  function b() {  
    if (n == 5) {  
      calls.add("b->c");  
      c();  
    }  
  }  
  
  function c() {  
    if (n == 5) {  
      calls.add("c->c");  
      n---;  
      c();  
    }  
  }  
  
  calls.add("main->a");  
  a();  
  
  console.log(calls);  
}
```



Different dynamic analyses, but with several commonalities

Specific runtime events to track

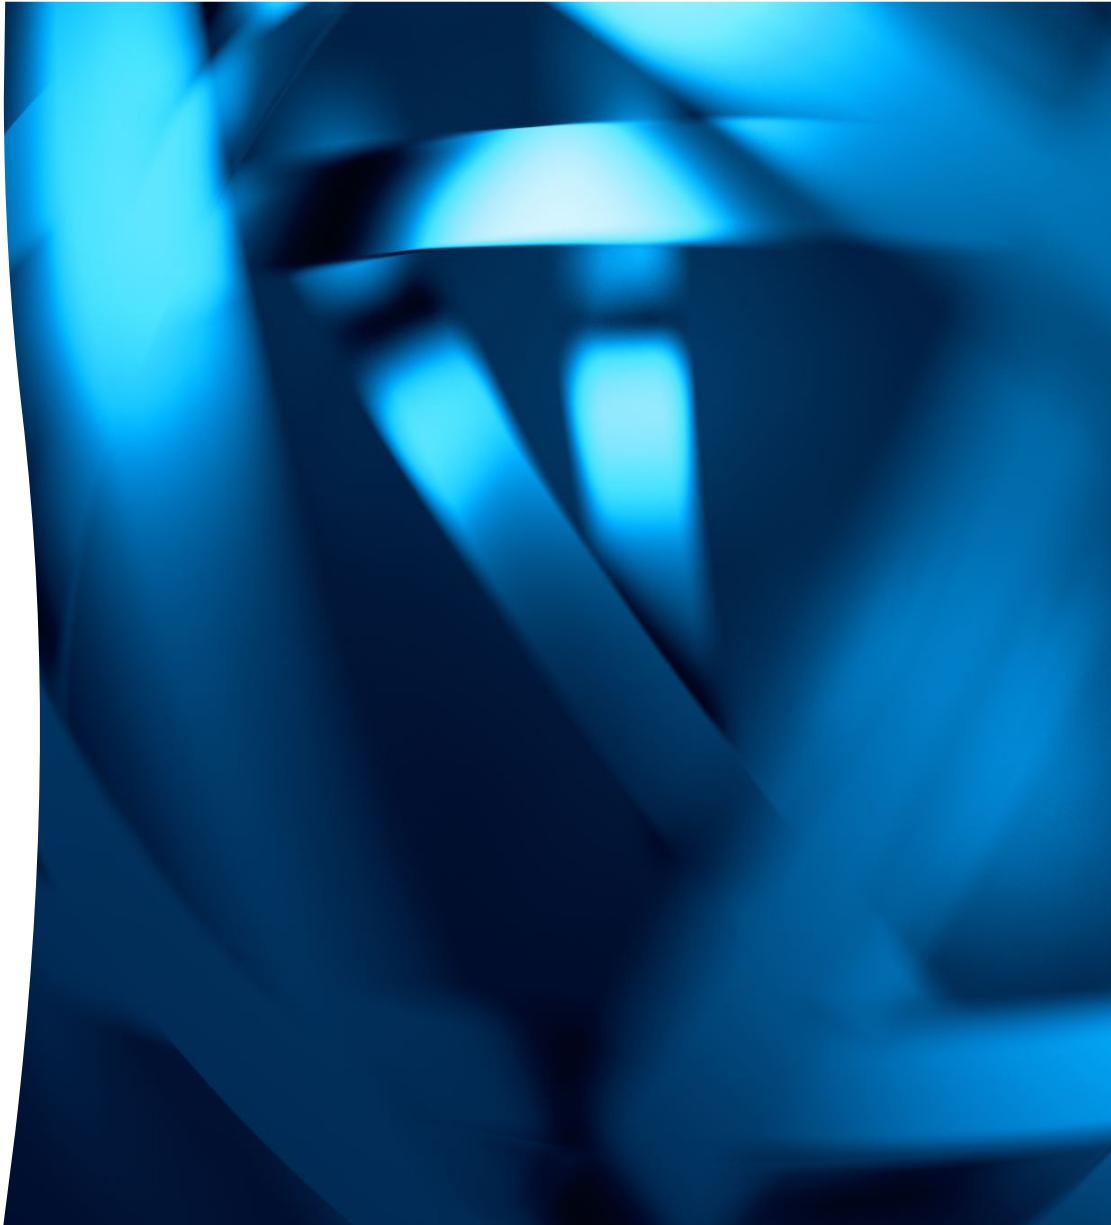
Analysis updates some state in response to events

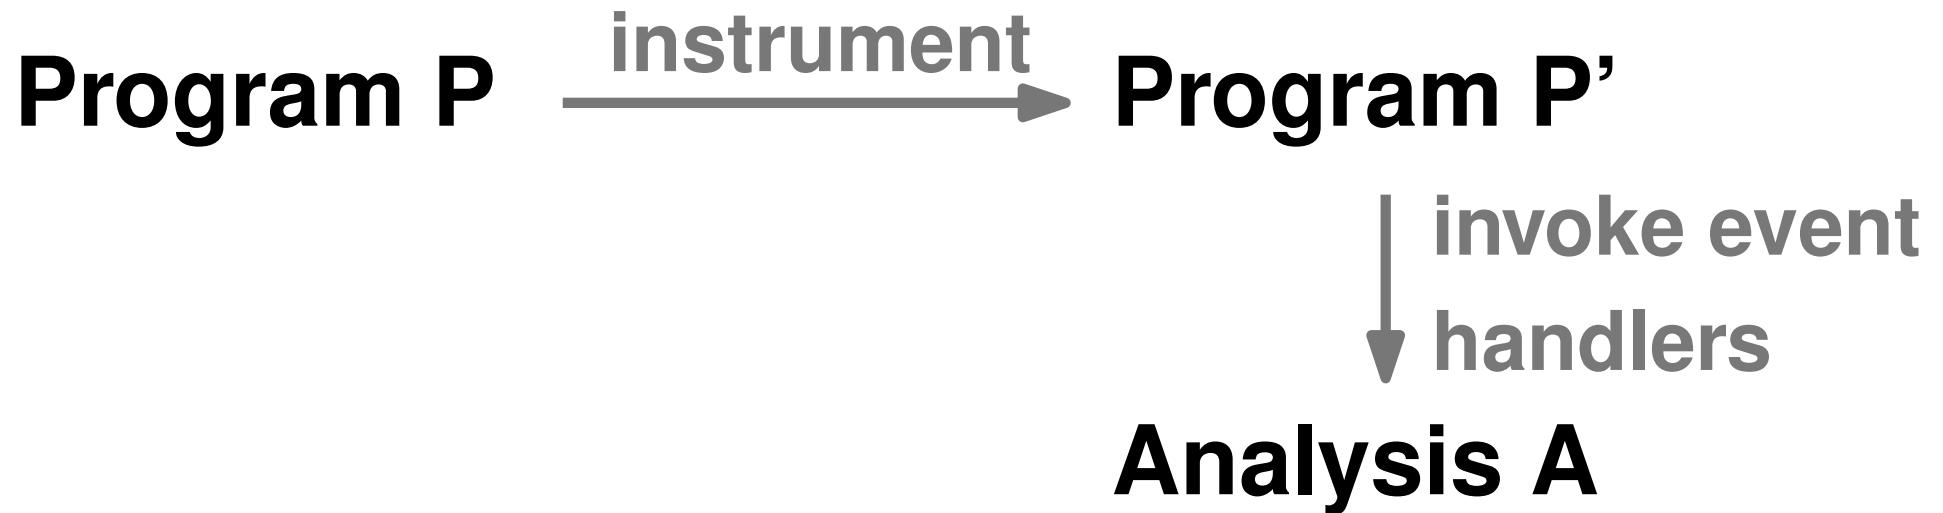
The ingredients

Identify the set (of kinds) of runtime events

The analysis can register for specific events

At runtime, instrumented program invokes event handlers





The idea (cont.)



Run time
events

Event	Example
Arithmetic operation	$2+3$
Boolean operation	<code>a > 0</code>
Branch	<code>if (c) ...</code>
Function call	<code>g()</code>
Return from function call	<code>x = g()</code>
Write into variable or field	<code>x.f = z</code>
Read of variable or field	<code>x.f = z</code>

Example

```
function main() {  
    let a = readInput();  
    let b = a + 3;  
  
    if (b === -23) {  
        foo();  
    } else {  
        b = 5;  
    }  
}
```

Example

```
function main() {  
    let a = readInput();  
    let b = a + 3;  
  
    if (b === -23) {  
        foo();  
    } else {  
        b = 5;  
    }  
}
```

Runtime events:

- Arithmetic operations
- Boolean operations
- Reads of variables
- Writes into variables
- Function calls

Input: -26

Example

```
function main() {  
    let a = readInput();  
    let b = a + 3;  
  
    if (b === -23) {  
        foo();  
    } else {  
        b = 5;  
    }  
}
```

Runtime events:

- Arithmetic operations
- Boolean operations
- Reads of variables
- Writes into variables
- Function calls

Input: -26

What sequence of events get triggered?

Example

```
function main() {  
    let a = readInput();  
    let b = a + 3;  
  
    if (b === -23) {  
        foo();  
    } else {  
        b = 5;  
    }  
}
```

Runtime events:

- Call of **readInput**
- Write of **-26** to **a**
- Read of **a** (-26)
- Arithmetic op (-26+3 = -23)
- Write of **-23** to **b**
- Read of **b** (-23)
- Boolean op (-23 == -23 aka true)
- Call **foo**

Input: -26

This is the sequence of events triggered!!

Remark and questions

Easy enough for small examples.

- But what happens with **large codebases**?
- How can we **formally guide the instrumentation of code and the analysis**?





Formally: Extended Operational Semantics

Tracking runtime events: Additional behavior performed during program execution

Formally describe by extending the operational semantics

Some Notation

Runtime events:

- Write of **-26** to **a** $\implies a \leftarrow -26$
- The if branch is taken \implies **if true**

Some Notation

CONFIGURATION

$$\langle P, s, e \rangle$$

P: program

s: store (mapping from program variables to values)

e: sequence of run-time events

Example

$$\frac{s(x) = n}{\langle !x, s \rangle \rightarrow \langle n, s \rangle} \text{ VAR}$$

Example

$$\frac{s(x) = n}{\langle !x, s \rangle \rightarrow \langle n, s \rangle} \text{ VAR}$$

$$\frac{s(x) = n}{\langle !x, s, e \rangle \rightarrow \langle n, s, e \rangle} \text{ VAR}$$



Run-time events are left unchanged

Example

$$\overline{\langle x ::= n, s \rangle \rightarrow \langle skip, s[x == n] \rangle} \quad \text{ASS}$$

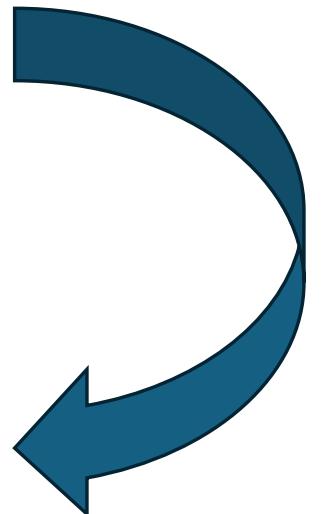
Example

$$\frac{}{\langle x ::= n, s \rangle \rightarrow \langle skip, s[x == n] \rangle}$$

ASS

$$\frac{}{\langle x ::= n, s, e \rangle \rightarrow \langle skip, s[x == n], e ; (x \leftarrow n) \rangle}$$

ASS



$$\overline{\langle \text{if true then } C1 \text{ else } C2, s, e \rangle \rightarrow \langle C1, s, e; (\text{if true}) \rangle}$$

QUIZ: Extend the operational rules for conditional – all cases

The Valgrind Framework

Valgrind is a *dynamic binary instrumentation framework*.

Valgrind runs programs inside a virtual machine and **monitors its behavior at runtime** instrumenting the source code.

The Instrumentation code

```
if (x === 3)  {  
    y = 3;  
    foo();  
    LogCall(foo);  
}
```

The instrumented AST

```
if (x === 3)  {  
    y = 3;  
    foo();  
    LogCall(foo);  
}
```

Program

```
└─ IfStatement  
    ├─ test: BinaryExpression (==)  
    │  ├─ left: Identifier("x")  
    │  └─ right: NumericLiteral(3)  
    └─ consequent: BlockStatement  
        ├─ ExpressionStatement  
        │  └─ AssignmentExpression (=)  
        │      ├─ left: Identifier("y")  
        │      └─ right: NumericLiteral(3)  
        └─ ExpressionStatement  
            └─ CallExpression  
                ├─ callee: Identifier("foo")  
                └─ arguments: []  
            ExpressionStatement  
            └─ CallExpression  
                ├─ callee: Identifier("LogCall")  
                └─ arguments:  
                    └─ Identifier("foo")  
    └─ alternate: null
```

Implementation steps

1. The instrumented AST **is compiled into** an intermediate form (IR).
2. The Valgrind IR allows handling of run-time events around memory operations, system calls, and thread interactions.
3. This IR runs on the Valgrind virtual machine
4. Valgrind intercepts and logs all the run-time events of interest.
5. When the program finishes (or fails), Valgrind produces a **detailed report** describing:
 - Where run-time events occurred (with call stacks)
 - What kind of run-time events they were
 - Which allocations caused memory leaks

Best Practice

A robust workflow for C/C++ (and analogous for other languages) often looks like this:

Static phase — run:

- clang-tidy: enforces code style, detects suspicious logic.
- cppcheck or Coverity: deeper semantic and data-flow analysis.

Dynamic phase — run:

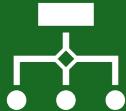
- valgrind --leak-check=full ./program: check runtime memory safety.
- valgrind --tool=helgrind ./program: check for race conditions.

What is the significance of all
this within the framework of this
course?

What have we learned?



Two key aspects:



Instrumenting the language and its operational semantics (understood as a guide to implementation) to capture and manage **run-time events**.



Analyzing run-time events to gain a better understanding of what **dynamic program analysis** means for the program under examination.



History dependent Access Control
A Running Example

History Dependent Access Control: Key Ideas

History-dependent access control: access permissions depend not only on *who* the principal is, but also on *what actions* have occurred previously in the program's execution.

Security policies as predicates over execution histories

access is granted only if a specific authentication step was executed earlier.

Provides a **formal model** connecting **program semantics** with **security enforcement**.

History dependent access control: discussion

History dependent access control allows one express access policies of the form:

Allow access only if this request has been authenticated by a trusted login program,

Allow a module to access resource only when invoked through certain routines



Ensures **fine-grained control: execution traces** are used to reason about **dynamic, context-sensitive access control**.

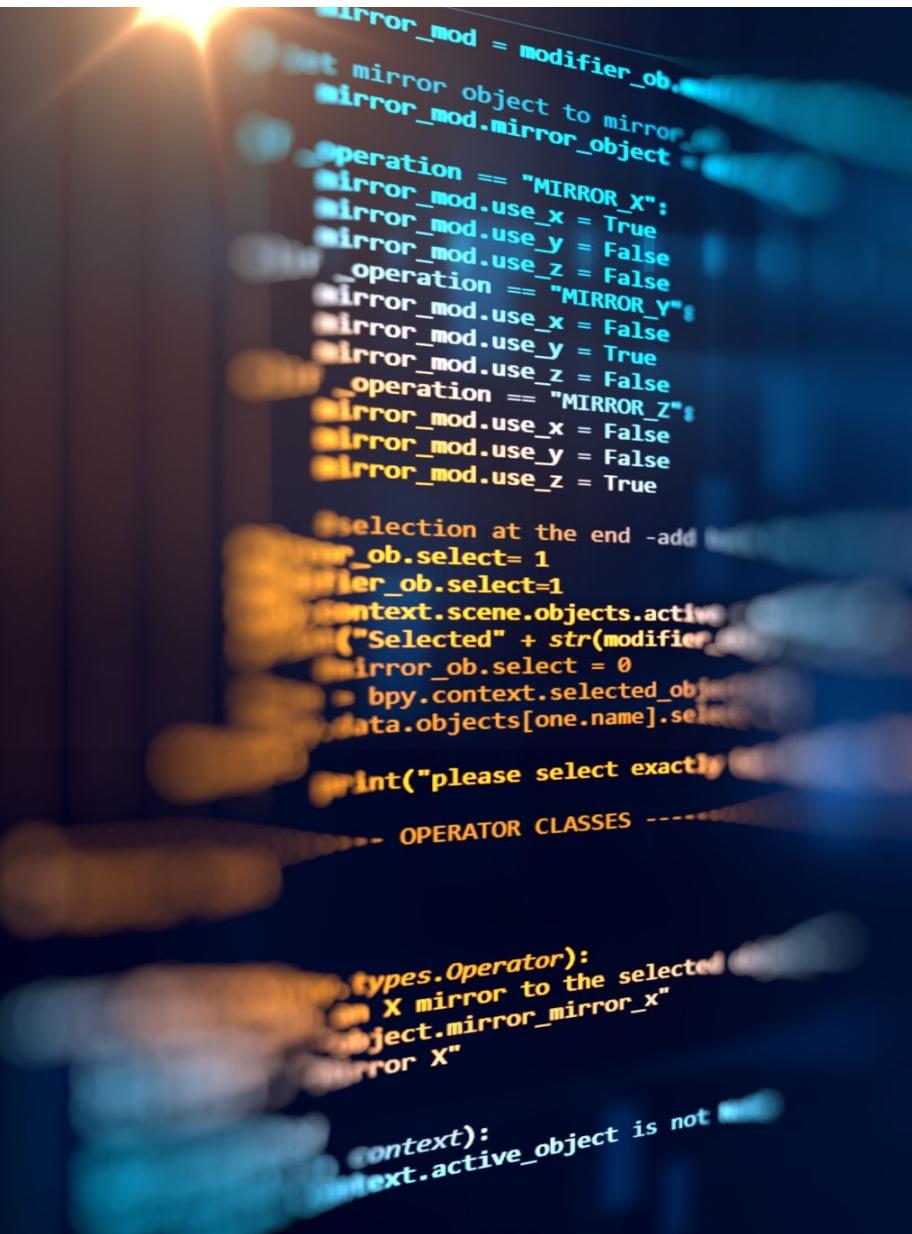
References

- **M. Abadi and C. Fournet.** *Access Control Based on Execution History.*
In *Proceedings of the 10th Annual Network and Distributed System Security Symposium (NDSS 2003)*, San Diego, CA, 2003.

References

M. Abadi and C. Fournet. Access Control Based on Execution History. In *Proceedings of the 10th Annual Network and Distributed System Security Symposium (NDSS 2003)*, 2003.

Massimo Bartoletti, Pierpaolo Degano, Gian-Luigi Ferrari: History-Based Access Control with Local Policies. In Proc, FoSSaCS 2005.



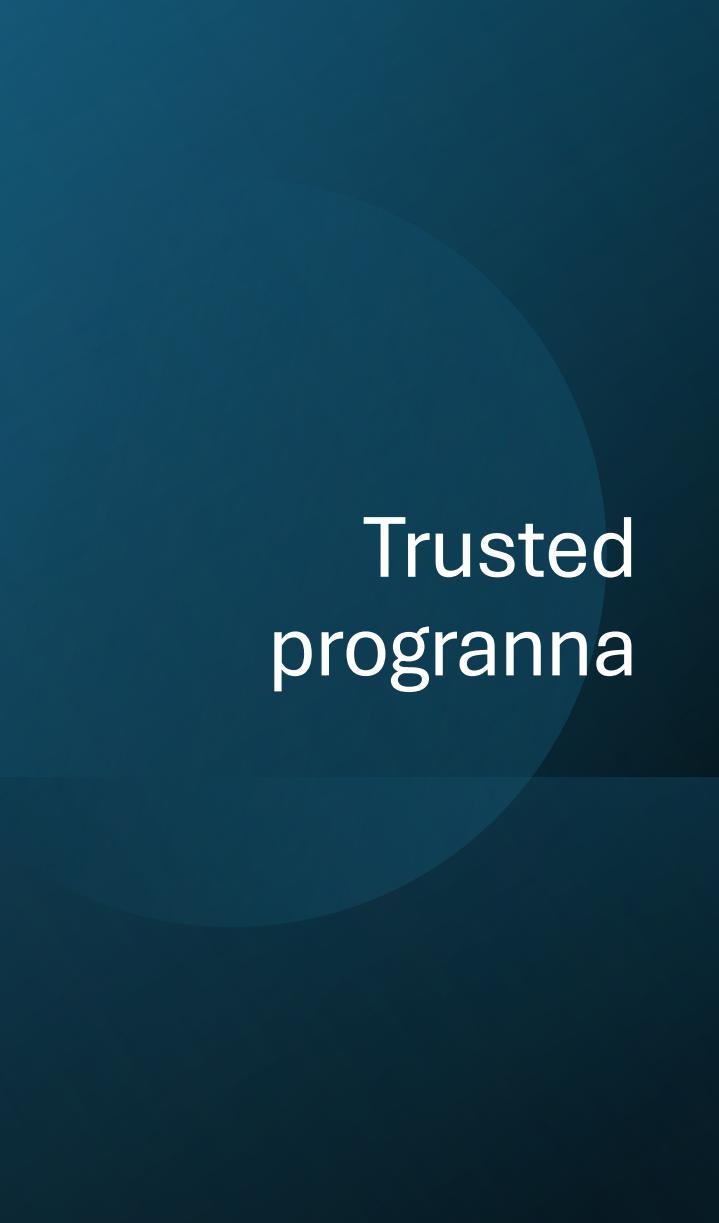
Motivating example

- Consider a simple **web browser** that displays HTML pages and executes suitable code.
- Code may be **trusted** (for instance, because they are signed or downloaded from a trusted third party) or **untrusted**.
- The browser enforces a **security policy** which is **always applied to untrusted code**:
 - The security policy states that code *cannot connect to the network after it has read from the local disk*.

We define the browser as a function that processes a URL u .
If the URL refers to an HTML page, the page is displayed.
If the URL refers to code, it is executed when trusted; otherwise, it is executed
under the control of a security policy.

```
(* u : url  *)
(* execute : (unit -> unit) -> unit *)
(* enforce : policy -> (unit -> unit) -> unit *)

let browser (u : url) : unit =
  if html u then
    display u
  else if trusted u then
    execute u
  else
    enforce p u  //p is the browser security policy
```



Trusted programma

We consider three trusted programs:

- **Read** to read files,
- **Write** to write files,
- **Connect** to open network connections.

These programs are overly simplified, because we are only interested in the events they execute.

Trusted Read

```
(* An action type *)
type action = Read | Write | Connect

(* Dispatcher: maps an action to a thunk (unit -> unit) *)
let action : action -> unit -> unit = function
  | Read  -> (fun () -> print_endline "Performing READ")
  | Write -> (fun () -> print_endline "Performing WRITE")
  | Connect -> (fun () -> print_endline "Performing CONNECT")

(* The function: ignores its argument and executes Read *)
let TrustedRead : 'a -> unit =
  fun _ -> action Read ()
```

Runtime enforcement: key idea

- **Execution monitors** enforce history-based security policies at run time
 - They **observe** computations and **abort** when a violation is imminent.
- **Events** = observations of **security-relevant actions**
 - Examples: opening a socket, reading/writing a file.
- **Histories** = (possibly infinite) **sequences of events**.
- **Policy as a global invariant**: must hold **at every point** during execution.
- **Expressive power**: execution monitors enforce **exactly the class of safety properties** (not liveness).



Trusted code

The behavior of trusted code, when executed within the browser, does not violate security policies;

for trusted code, the security monitor is vacuous (i.e., disabled).

Untrusted code

```
(* An action type *)
type action = Read | Write | Connect

(* Dispatcher: maps an action to a thunk (unit -> unit) *)
let action : action -> unit -> unit = function
  | Read  -> (fun () -> print_endline "Performing READ")
  | Write -> (fun () -> print_endline "Performing WRITE")
  | Connect -> (fun () -> print_endline "Performing CONNECT")

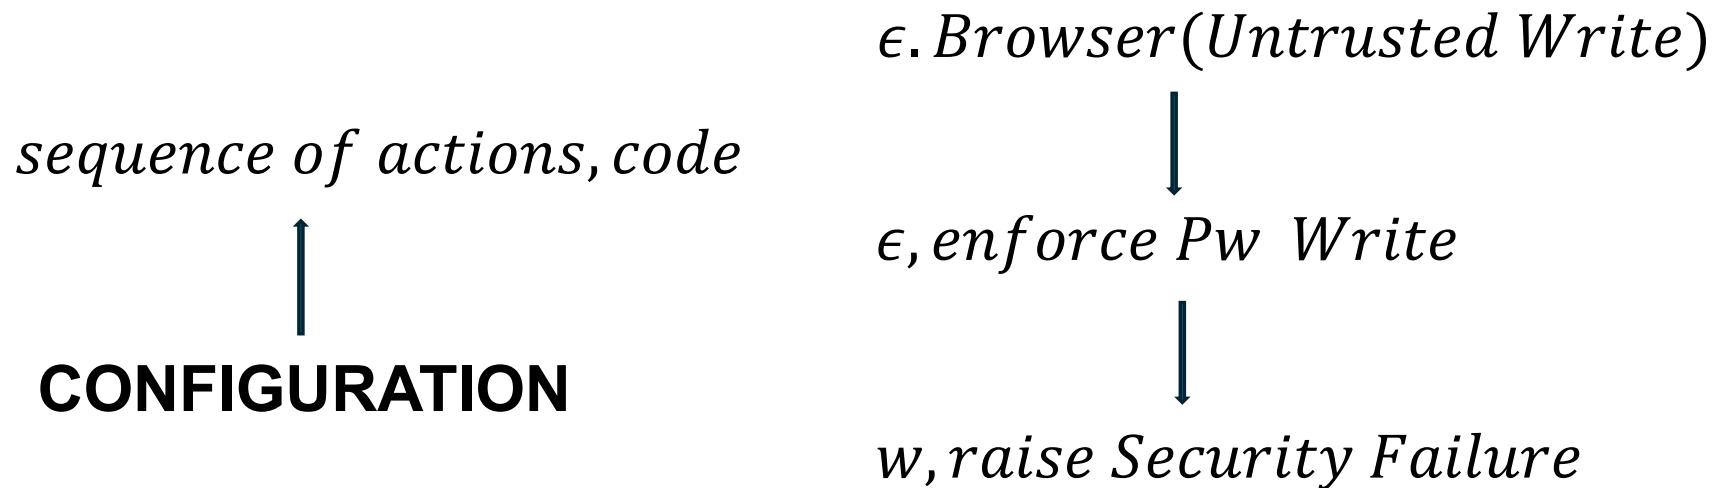
(* The function: idoes not gnore its argument *)
let Untrusted : url ->'a -> unit =
  fun (z:url) _ -> Browser z
```

Execution

What is the behaviour of an untrusted code (**Untrusted Write**) executed under the security policy **Pw** stating that untrusted code cannot write the local file system?

Execution

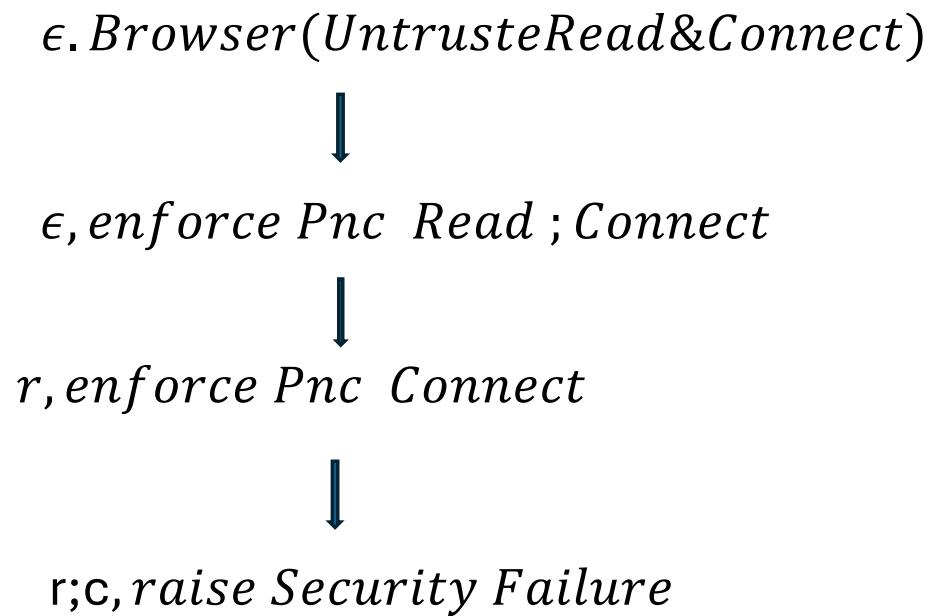
The behaviour of an untrusted code (**Untrusted Write**) executed under the security policy **Pw** stating that untrusted code cannot write the local file system, is illustrated by the following trace:



Execution

The behaviour of an untrusted code (**Untrusted Read&Connect**) executed under the security policy **Pnc** stating that untrusted code cannot connect to the network after reading the local file system, is illustrated by the following trace:

**we have a security exception,
because the r;c does not
satisfy the security policy Pnc**



The Formal Representation



Syntax

EXPRESSIONS: Functional programming language

$$e, e' ::= x \mid \alpha \mid \text{if } b \text{ then } e \text{ else } e' \mid \lambda_z x. e \mid e e'$$

PROGRAMS: Sandbox monitoring execution under the policy

$$P ::= \varphi[e]$$

Execution Monitors

Runtime Monitoring: Checking policy adherence during execution (e.g., sandboxing)

Enforcement Mechanisms in VM: Virtual machines (e.g., WebAssembly, JVM) enforce security through restricted execution environments.

Execution Monitors: Intercepts security-sensitive operations and ensures they comply with policies.

Some challenging questions

Can we prove that mechanism **M** enforces the security policy **P**?

- What is the mathematical definition of a policy?
- What is the programming abstractions for declaring security policies?
- What does it mean to enforce a policy within a programming language?

Are there limits to what is enforceable?

- Which enforcement approaches are best suited to which

Policies?

- Are there some policies that are completely beyond any known enforcement strategy?
- Are some enforcement approaches strictly more powerful than others?

OVERALL

- what is the landscape of policies, policy classes, and enforcement mechanisms?

Enforceable Security Policies

Enforceable Security Policies

F. Schneider, [ACM Transactions on Information and System Security](#), February 2000

Abstract

A precise characterization is given for the class of security policies enforceable with mechanisms that work by monitoring system execution, and automata are introduced for specifying exactly that class of security policies. Techniques to enforce security policies specified by such automata are also discussed.

The Ideal Execution Monitor

1. Sees *everything* a program is about to do before it does it
2. Can *instantly* and *completely* stop program execution (or prevent action)
3. Has *no other effect* on the program or system

The Ideal Execution Monitor

1. **Sees *everything* a program is about to do before it does it**
2. **Can *instantly* and *completely* stop program execution (or prevent action)**
3. **Has *no other effect* on the program or system**

**Can we implement this?
Probably not**

~~Real~~ Ideal Execution Monitor

1. ~~Sees everything a program is about to do before it does it~~
most things
2. ~~Can instantly and completely stop program execution (or prevent action)~~
3. ~~Has no other effect on the program or system~~
limited

Execution Monitor

Execution Monitors (EMs)

- EMs watch untrusted programs at runtime
- Events (raised by the run time executions) are mediated by the EM
- Violations solicit EM interventions (e.g. termination)

Example: File system access control

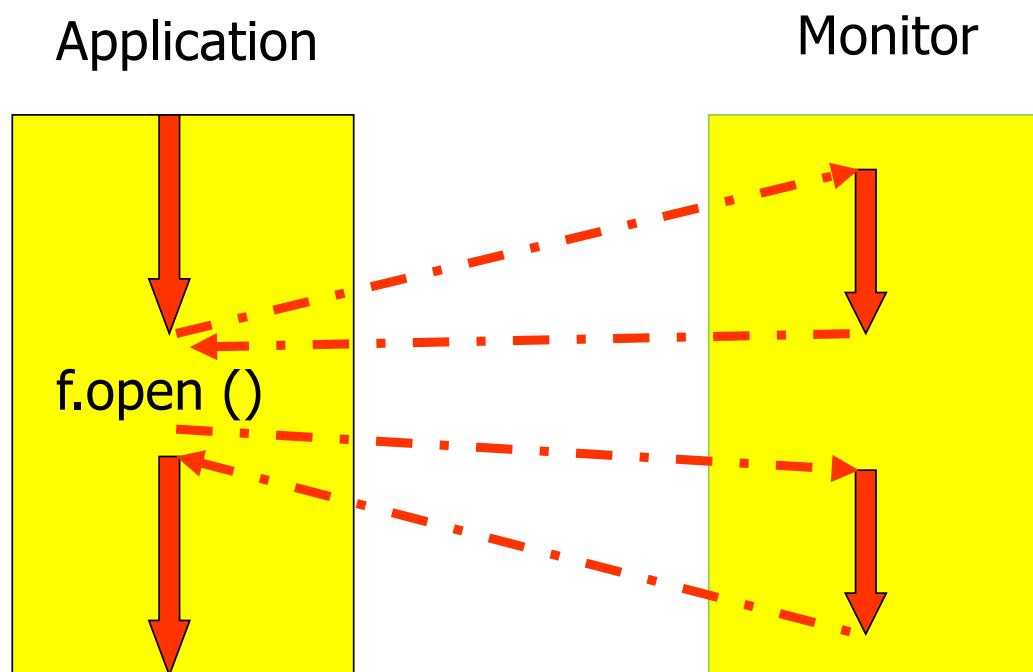
- EM is inside the OS
- decides policy violations using access control lists (ACLs)



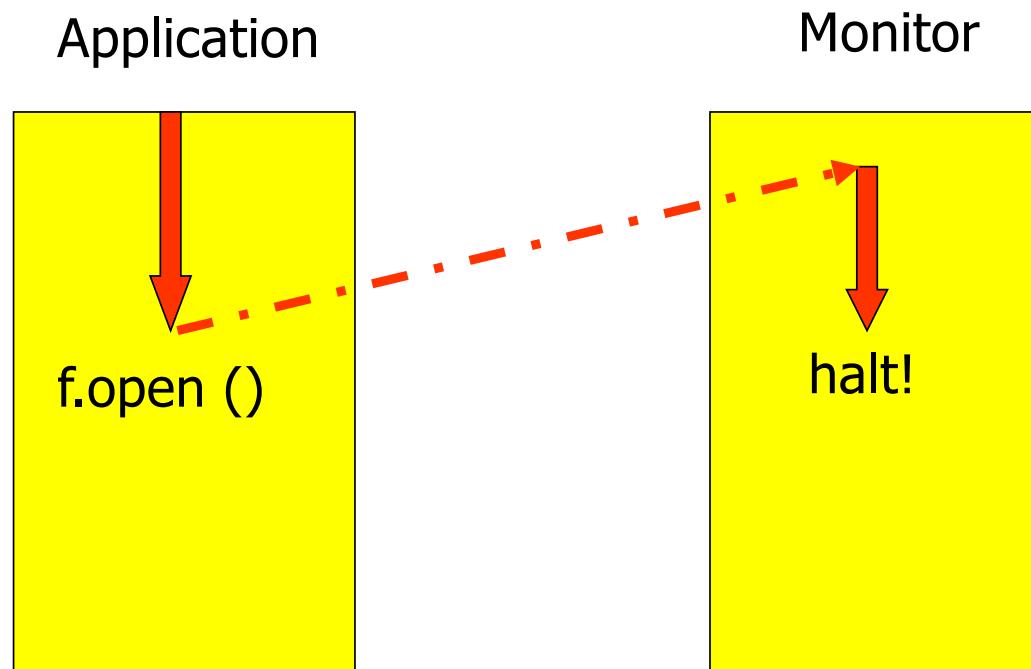
Execution Monitor

- EMs are run-time modules that runs in parallel with an application
 - Tracks execution flow at runtime to **detect and prevent** security violations dynamically.
 - monitor decisions may be based on **execution history**

EM: Good Operations



EM: Bad Operation



Execution and Security Policies

What's a program (at run-time)?

- A set of possible executions

What's an execution?

- A sequence of states

What's a security policy?

- A predicate on a set of executions

Definitions and Notations

- An *execution (or trace)* s is a sequence of security-relevant program events e (also called actions)
- Sequence may be **finite** or **(countably) infinite**
 - $s = e_1; e_2; \dots; e_k; e_{\text{halt}}$
 - $s = e_1; e_2; \dots; e_k; \dots$
 - The empty sequence ε is an execution
 - If s is the execution $e_1; e_2; \dots; e_i; \dots; e_l; \dots$ then $s[i]$ is the execution $e_1; e_2; \dots; e_i;$
- We simplify the formalism.
 - We model program termination as an infinite repetition of e_{halt} event.
 - Result: now all executions are infinite length sequences

Definitions and Notations (cont.)

- A program **S** is a **set** of sequences (possible executions)
 - A program is modelled as the set $S = \{s_1, s_2, \dots\}$
- A security policy **P** is a **property** of programs
- A policy partitions the program space into two groups:
 - Permissible
 - Impermissible
- Impermissible programs are censored somehow (e.g., terminated on violating runs)

Execution Monitors cannot enforce all Security Policies

- Some policies depend on:
 - Knowing about the future
 - If the program charges the credit card, it must eventually ship the goods
 - Knowing about all possible executions
 - Information flow – can't tell if a program reveals secret information without knowing about other possible executions
- Execution Monitors can only know about past of this particular execution

EM-enforceable policies

1. EM policies are universally quantified predicates over executions
 - $\forall s. P(s)$
 - **Policy P is called the detector.**

EM-enforceable policies

1. EM policies are universally quantified predicates over executions
 - $\forall s. P(s)$
 - **Policy P is called the detector.**
2. The detector must be prefix-closed
 - $P(\varepsilon)$ holds
 - $P(s;e)$ holds then $P(s)$ holds

EM-enforceable policies

1. EM policies are universally quantified predicates over executions
 - $\forall s. P(s)$
 - **Policy P is called the detector.**
2. The detector must be prefix-closed
 - $P(\varepsilon)$ holds
 - $P(s;e)$ holds then $P(s)$ holds
3. If the detector is not satisfied by a sequence (the detector rejects the sequence) then it must do so in finite time
 - $\neg P(s) \Rightarrow \exists i \neg Ps[i]$

EM-enforceable policies

1. EM policies are universally quantified predicates over executions
2. The detector must be prefix-closed
3. If the detector is not satisfied by a sequence (the detector rejects the sequence) then it must do so in finite time

4. Fact

- A policy satisfies (1), (2), and (3) if and only if it is a *safety policy*
- Lamport 1977: Safety policies say that some “bad thing” never happens
- EMs enforce safety policies!

Safety properties: Nothing bad ever happens

Safety property can be enforced **using only traces** of program

- If $P(t)$ does not hold, then all extensions of t are also bad

Amenable to **run-time enforcement**: don't need to know future

Examples:

- **access control** (e.g. checking file permissions on file open)
- **memory safety** (process does not read/write outside its own memory space)
- **type safety** (data accessed in accordance with type)

Formally...

- Ψ : set of all possible executions (can be infinite)
- Σ_S : set of executions possible by target program S
- P : security policy
 - set of executions \rightarrow Boolean

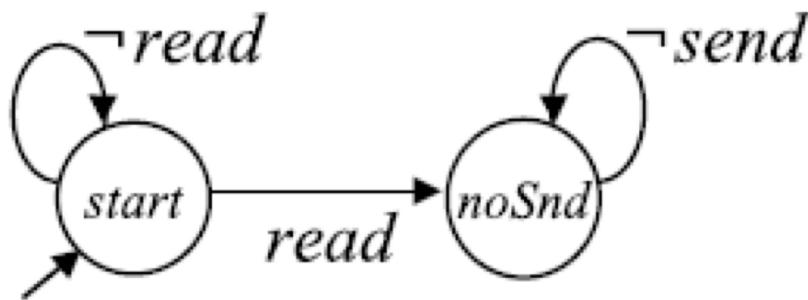
Program S is secure wrt the security policy P iff $P(\Sigma_S)$ is true.



Security Automata

- Security Automata (Erlingsson & Schneider 1999)
- Formalization of security policies
 - finite state automaton
 - accepts language of permissible executions
 - alphabet = set of events
 - edge labels = event predicates
 - all states accepting (language is prefix-closed)

Security Automata (Example)



NO SENDS AFTER READS

Execution Monitor

```
type event = Read of string | Write of string | Connect of string
type history = event list
type policy = history -> event -> bool (* true = allowed *)

type caps = {
  read  : string -> string;
  write : string -> unit;
  connect : string -> unit;
}

let sandbox (phi : policy) (base : caps) (prog : caps -> 'a) : 'a =
  let hist = ref [] in
  let check e =
    if phi !hist e then hist := e :: !hist
    else failwith "Policy violation"
  in
  :
}
```

Remark

- **What “global security policy” means here**
- The policy ϕ is an **invariant** over **all prefixes** of the event trace (a safety property).
- Every attempted effect is checked **before** it happens; violations abort execution.

Syntax

EXPRESSIONS

$$e, e' ::= x \mid \alpha \mid \text{if } b \text{ then } e \text{ else } e' \mid \lambda_z x. e \mid e e'$$

PROGRAMS

$$P ::= \varphi[e]$$

Operational semantics

- The configurations are pairs
$$\eta, e$$
- A transition $\eta, e \rightarrow \eta', e'$ indicates that, starting from a history η , the expression e may evolve to e' , possibly extending the history η .
- We write $\eta \models \varphi$ when the history η satisfies the policy φ .

Operational semantics

$$\frac{\eta, e_1 \rightarrow \eta', e'_1}{\eta, e_1 e_2 \rightarrow \eta', e'_1 e_2}$$

$$\frac{\eta, e_2 \rightarrow \eta', e'_2}{\eta, v e_2 \rightarrow \eta', v e'_2}$$

$$\frac{}{\eta, (\lambda_z x. e)v \rightarrow \eta, e\{v/x, \lambda_z x. e/z\}}$$

Operational Semantics

$$\mathcal{B}(b) = \text{true}$$

$$\eta, \text{if } b \text{ then } e_0 \text{ else } e_1 \rightarrow \eta, e_0$$

$$\mathcal{B}(b) = \text{false}$$

$$\eta, \text{if } b \text{ then } e_0 \text{ else } e_1 \rightarrow \eta, e_1$$

Operational Semantics

$$\eta, \alpha \rightarrow \eta\alpha, *$$

Operational Semantics

$$\frac{\eta, e \rightarrow \eta', e' \quad \eta, \eta' \models \varphi}{\eta, \varphi[e] \rightarrow \eta', \varphi[e']}$$

$$\frac{\eta \models \varphi}{\eta, \varphi[v] \rightarrow \eta, v}$$



Summary

- **Dynamic Analysis of History-Dependent Access Control**
- We have shown how to **characterize the dynamic analysis of history-dependent access control**, defining its **formal structure** and its **limitations** (safety security policies).

Next step

Can we do better?
Can we **integrate static**
and dynamic analysis?



Next lecture