

EXERCISE #25

LINTING REVIEW

Write your name and answer the following on a piece of paper


Give an example of a legal program in C that a linter would nevertheless flag



Quiz 2 Grades forthcoming

Quiz 2 recap lecture

**ADMINISTRIVIA
AND
ANNOUNCEMENTS**



PROGRAM INSTRUMENTATION

EECS 677: Software Security Evaluation

Drew Davidson

LAST TIME: LINTING

REVIEW: LAST LECTURE


REMOVING THE “STRAY FIBERS” OF A PROGRAM


Analyze common “anti-patterns” that are likely to cause issues (security and otherwise)

NOTABLE ANALYSIS TOOLS

Lint – The original analysis tool

Splint – Security analysis tool





CLASS PROGRESS

MOVING ON TO DYNAMIC ANALYSIS


More heuristic by nature

LIMITS OF STATIC ANALYSIS

PROGRAM INSTRUMENTATION: BASIC IDEA

PRACTICAL ISSUES

- Unsoundness of bug finding / incompleteness of program verification
- Scalability
- Significant engineering effort



REVISING DYNAMIC ANALYSIS

PROGRAM INSTRUMENTATION: BASIC IDEA

GIVING UP ON GUARANTEES

- Finding bugs (even “low-hanging fruit”) is useful!

ADVANTAGES

- Simplest form: testing
- Scalability ~~X~~ - Partial credit



BEYOND TESTING

PROGRAM INSTRUMENTATION: BASIC IDEA

LIMITATIONS OF “PLAIN” TESTING

- Property may not be immediately observable from output alone
- The circumstances under which the issue occurs may not be obvious



INSERTING PROGRAM PROBES

PROGRAM INSTRUMENTATION: BASIC IDEA

INSERT CHECKS / REPORTS INTO THE ANALYSIS TARGET

Addresses both of the previous issues – can report upon program state and even program path

A NEW CONCERN – THE EFFICIENCY OF THE (INSTRUMENTED) PROGRAM

Potential slowdown on each program path

LACK OF HOLISTIC INFORMATION

Somewhat limited by the information the probes can report

what does running a probe gain vs cost



what does calculate smart probe placement gain vs cost?



EXAMPLE: CONTROL PROFILING

PROGRAM INSTRUMENTATION: BASIC IDEA

COUNTING HOW MANY TIMES CERTAIN BEHAVIORS OF THE PROGRAM ARE EXERCISED

Why is this useful? (Placing sanitizers)

THIS ACTUALLY TURNS OUT TO BE A LITTLE BIT TRICKY!

Actually turns out to be a little bit tricky!

We'll describe some of the issues / solution as per Ball and Larus, '96

Ball and Larus '96



```

V = SECRET
for (i = 1; i < 1000000) {
  i++;
}
LEAK(V);
  
```

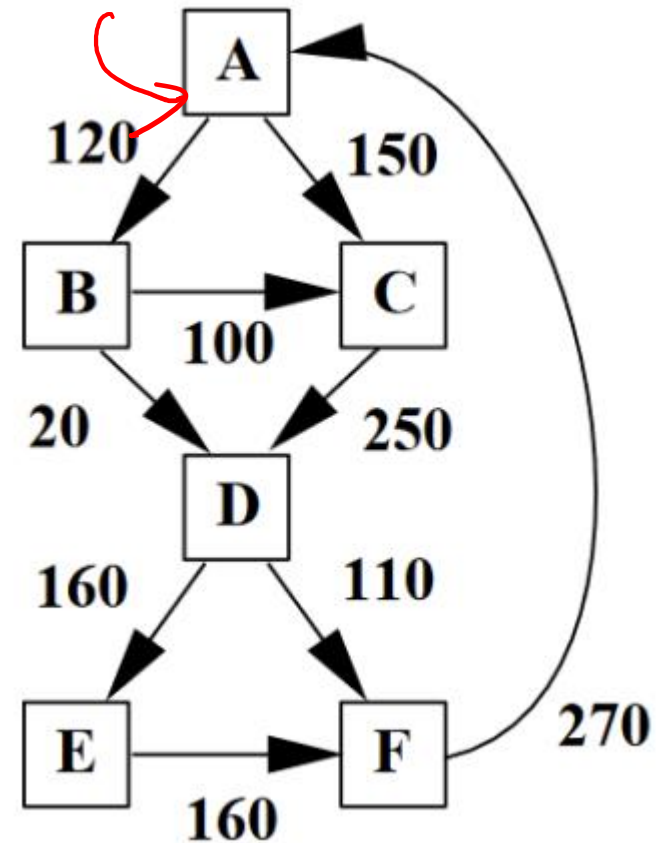
BRANCH FREQUENCY

PROGRAM INSTRUMENTATION: APPROACH

NAÏVE APPROACH: INSTRUMENT
PROBES AT EACH EDGE

Inefficient!

We don't really need an A → B counter
(it's the sum of the B → C and B → D counters)



EXAMPLE: COVERAGE / FREQUENCY

PROGRAM INSTRUMENTATION: APPROACH

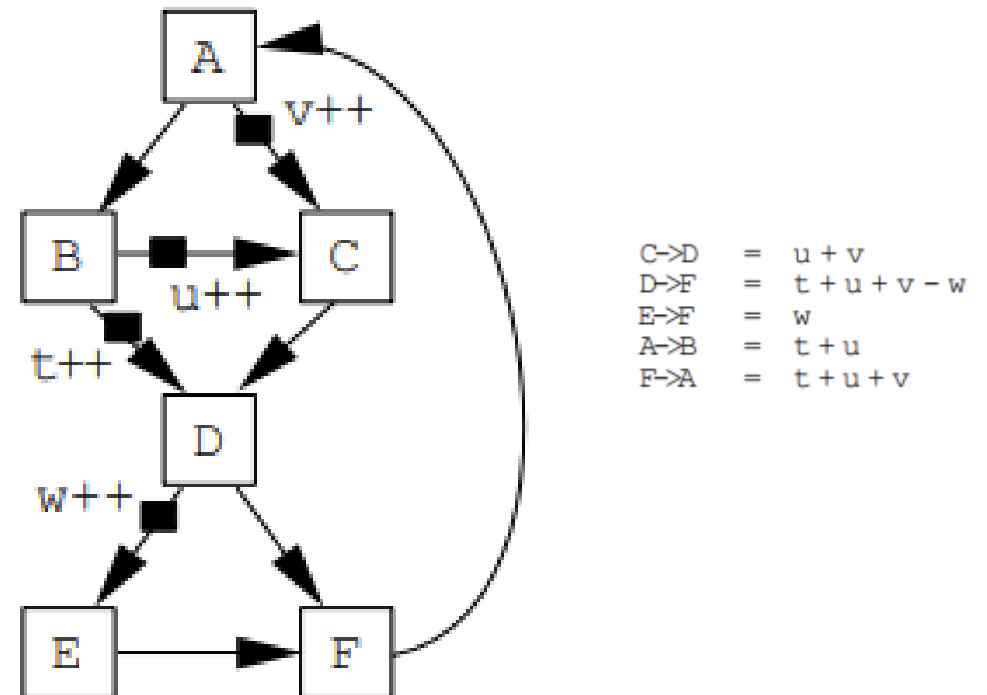
EXAMPLE OF INSTRUMENTATION: COUNTING EXECUTION FREQUENCY

Why is this useful? (Placing sanitizers)

Let's first consider inserting edge counters

Inefficient!

We don't really need an A -> B counter
(it's the sum of the B -> C and B -> D counters)




PATH FREQUENCY

PROGRAM INSTRUMENTATION: APPROACH

NAÏVE IMPLEMENTATION: SUM UP
EDGE COUNTERS

Path	Prof1	Prof2
ACDF	90	110
ACDEF	60	40
ABCDF	0	0
ABCDEF	100	100
ABDF	20	0
ABDEF	0	20



EFFICIENT PATH AND BRANCH COUNTERS

PROGRAM INSTRUMENTATION: APPROACH

Efficient Path Profiling

Thomas Ball
Bell Laboratories
Lucent Technologies
tball@research.bell-labs.com

James R. Larus*
Dept. of Computer Sciences
University of Wisconsin-Madison
larus@cs.wisc.edu

BALL AND LARUS '96

Intuition:

- Assign integer values to edges such that no two paths compute the same path sum (Section 3.2).
- Select edges to instrument using a spanning tree

Abstract

A path profile determines how many times each acyclic path in a routine executes. This type of profiling subsumes the more common basic block and edge profiling, which only approximate path frequencies. Path profiles have many potential uses in program performance tuning, profile-directed compilation, and software test coverage.

This paper describes a new algorithm for path profiling. This simple, fast algorithm selects and places profile instrumentation to minimize run-time overhead. Instrumented programs run with overhead comparable to the best previous profiling techniques. On the SPEC95 benchmarks, path profiling overhead averaged 31%, as compared to 16% for efficient edge profiling. Path profiling also identifies longer paths than a previous technique, which predicted paths from edge profiles (average of 88, versus 34 instructions). Moreover, profiling shows that the SPEC95 train input datasets covered most of the paths executed in the test datasets.




Figure 1. Example in which edge profiling does not identify the most frequently executed paths. The table contains two different path profiles. Both path profiles induce the same edge execution frequencies, shown by the edge frequencies in the control-flow graph. In path profile *Prof1*, path *ABCDE F* is most frequently executed, although the heuristic of following edges with the highest frequency identifies path *ACDEF* as the most frequent.

1 Introduction

Program profiling counts occurrences of an event during a program's execution. Typically, the measured event is the execution of a local portion of a program, such as a routine or line of code. Recently, fine-grain profiles—of basic blocks and control-flow edges—have become the basis for profile-driven compilation, which uses measured frequencies to guide compilation and optimization.

*This research supported by: Wright Laboratory Avionics Directorate, Air Force Material Command, USAF, under grant #F33615-94-1-1525 and ARPA order no. B550; NSF NYI Award CCR-9357779, with support from Hewlett Packard, Sun Microsystems, and PGI; NSF Grant MIP-9225097; and DOE Grant DE-FG02-93ER25176. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Wright Laboratory Avionics Directorate or the U. S. Government.

One use of profile information is to identify heavily executed paths (or traces) in a program [Fis81, Ell85, Cha88, YS94]. Unfortunately, basic block and edge profiles, although inexpensive and widely available, do not always correctly predict frequencies of overlapping paths. Consider, for example, the control-flow graph (CFG) in Figure 1. Each edge in the CFG is labeled with its frequency, which normally results from dynamic profiling, but in the figure is induced by *both* path profiles in the table. A commonly used heuristic to select a heavily executed path follows the most frequently executed edge out of a basic block [Cha88], which identifies path *ACDEF*. However, in path profile *Prof1*, this path executed only 60 times, as compared to 90 times for path *ACDF* and 100 times for path *ABCDE F*. In profile *Prof2*, the disparity is even greater although the edge profile is exactly the same.

This inaccuracy is usually ignored, under the assumption that accurate path profiling must be far more expensive than basic block or edge profiling. Path profiling is the ultimate form of control-flow profiling, as it uniquely deter-

INSTRUMENTATION APPROACHES

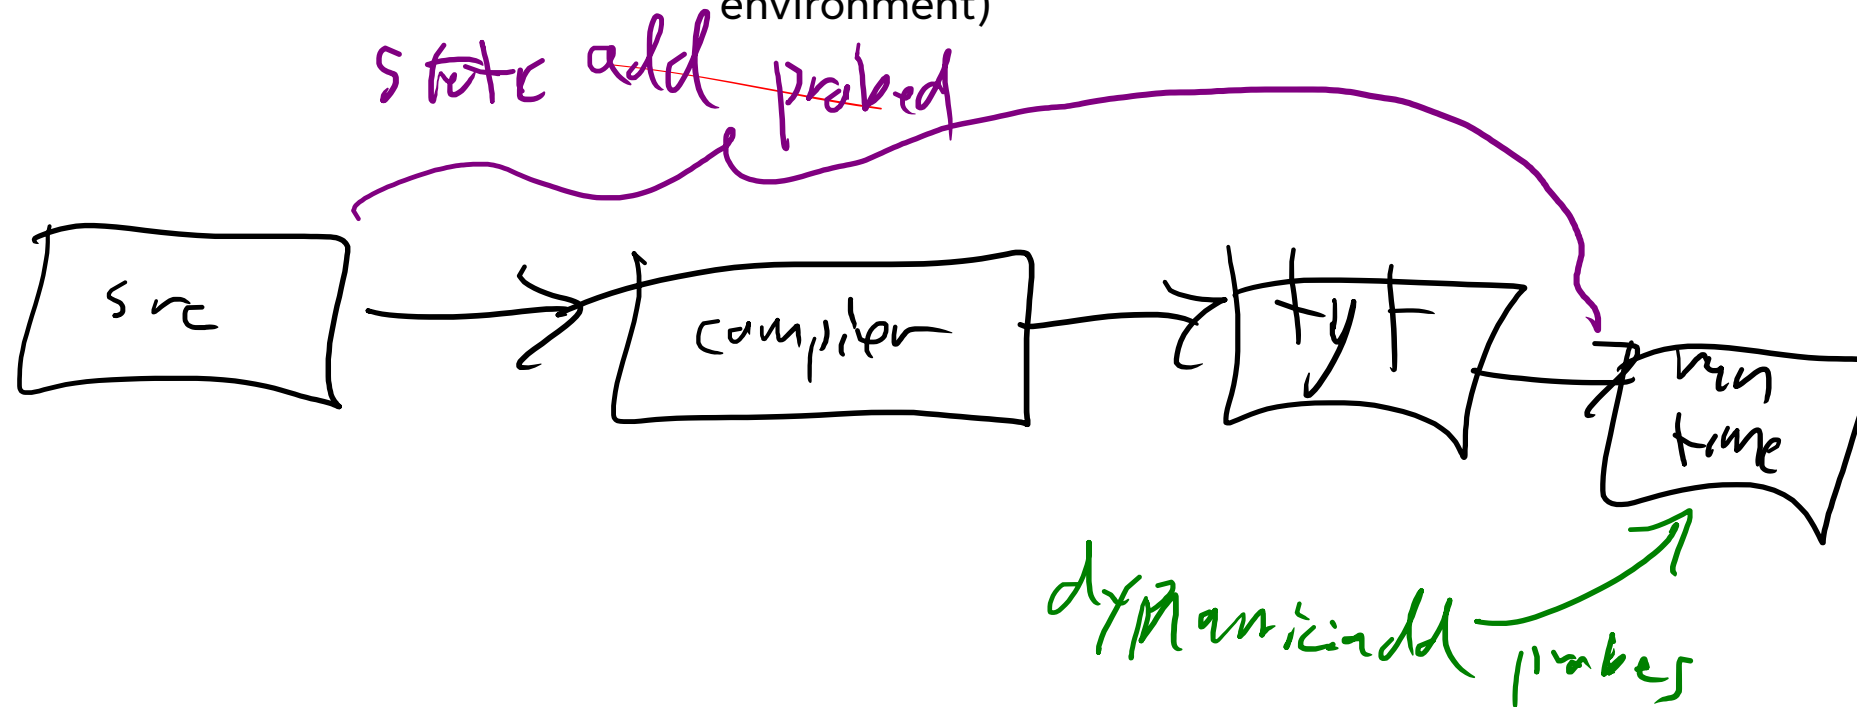
PROGRAM INSTRUMENTATION: APPROACH

STATIC INSTRUMENTATION

Add probes before the program is run (i.e. rewrite the program executable)

DYNAMIC INSTRUMENTATION

Probe while the program is run (i.e. insert probes just-in-time or as part of the environment)



DYNAMIC INSTRUMENTATION TOOLS

PROGRAM INSTRUMENTATION: APPROACH

FREQUENTLY INVOLVE A CUSTOM RUNTIME

Add probes before the program is run (i.e.
rewrite the program executable)

EXAMPLES

Intel PIN

GDB

DYNAMIC INSTRUMENTATION EXAMPLE: GDB

PROGRAM INSTRUMENTATION: APPROACH

```
1 #include "stdio.h"
2
3 int main(){
4     int a = 1;
5     if (a == 1){
6         printf("TRUE BRANCH: A is %d\n", a);
7     } else {
8         printf("FALSE BRANCH: A is %d\n", a);
9     }
10 }
```

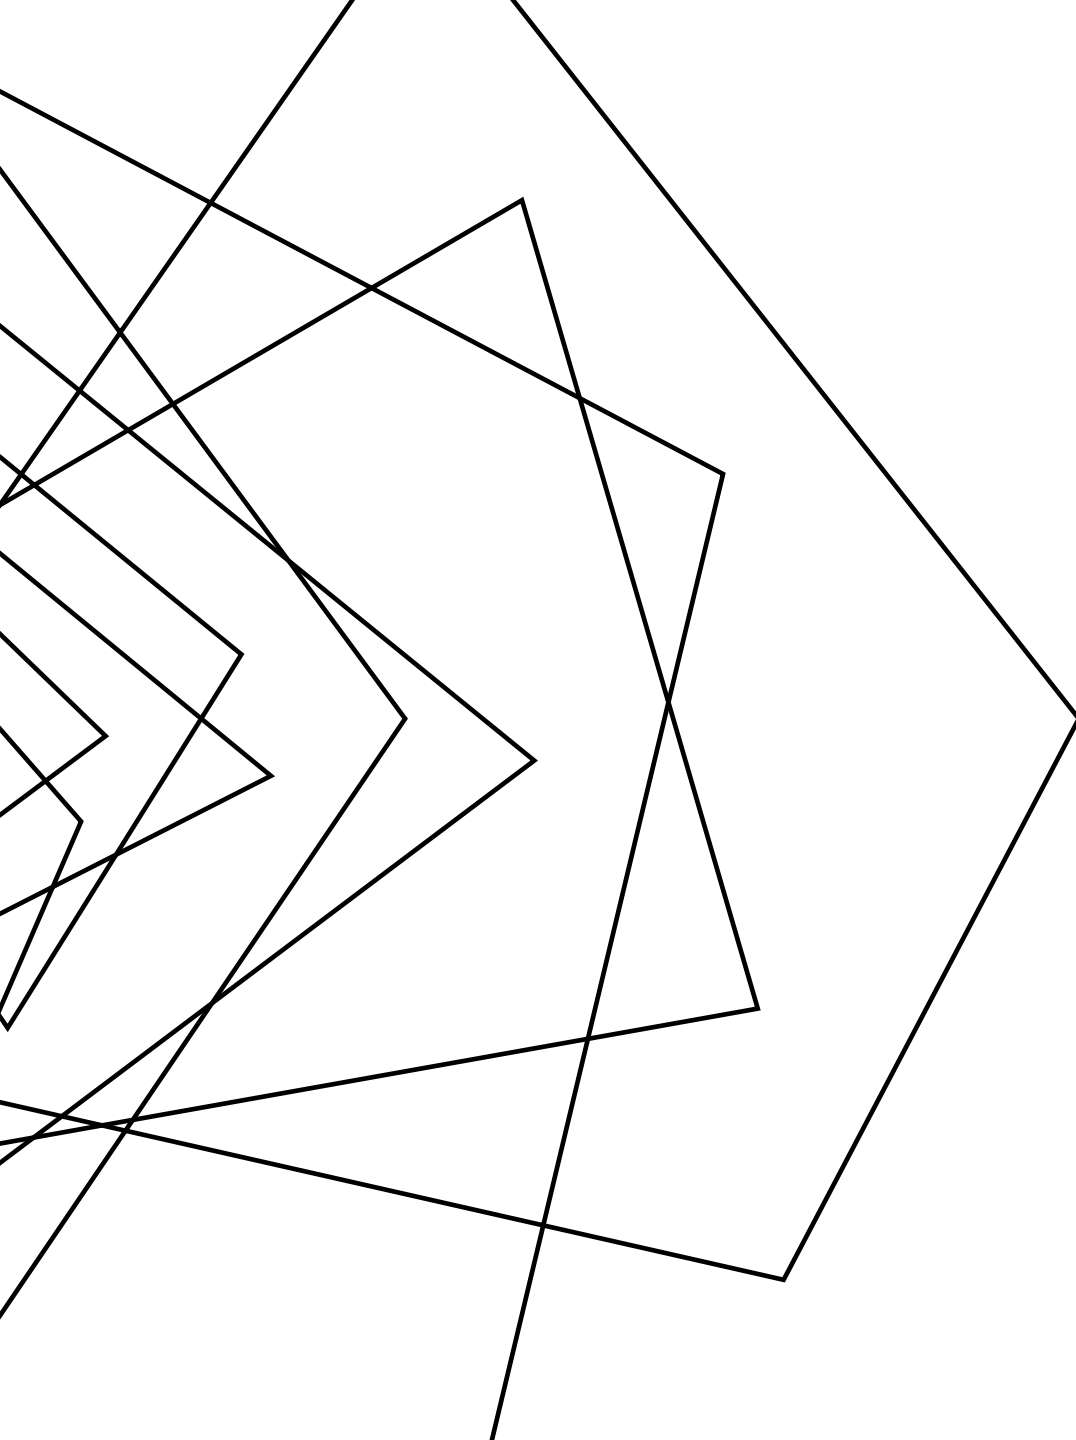
gcc -O0 -g prog.c -o prog

gdb prog

b 5

set variable a = 3

n



WRAP-UP

WE'VE DESCRIBED 2 FORMS OF
ALTERING THE PROGRAM

More heuristic by nature