

## **Class 18**

- Questions/comments
- Discussion of academic honesty, GT Honor Code
- Efficient path profiling
- Final project presentations: Dec 1, 3; 4:35-6:45
- Assign (see Schedule for links)
  - Problem Set 7 discuss
  - Readings

1

## **Execution Tracing and Profiling**

- Gathering dynamic information about programs
  - Execution coverage
  - Execution profiling
  - Execution tracing

2

## Execution Tracing and Profiling

- Gathering dynamic information about programs
  - Execution coverage
  - Execution profiling
  - Execution tracing
- Three main alternatives
  - Debugging interfaces
  - Customized runtime systems
  - Instrumentation
    - Post-processing
    - Online processing
    - Preprocessing

3

## Debugging Interfaces

- **Debugging Interfaces** provide hooks into the runtime system that allow for collecting various dynamic information while the program executes.
- Examples:
  - Java Platform Debugger Architecture (JPDA)
    - JVM Debugging Interface (JVMDI)
    - JVM Profiling Interface (JVMPPI)
    - Java Virtual Machine Tool Interface (JVMTI) [New]
  - Valgrind
  - DynamoRIO
  - Emulators for embedded systems

4

## Customized Runtime Systems

- **Customized Runtime Systems** are runtime systems modified to collect some specific dynamic information.
- Examples:
  - Jalapeño JVM

5

## Instrumentation Tools

- Source-level
  - EDG parser (AST)
  - Customized gcc
- Binary/bytecode level
  - Vulcan
  - BCEL
  - SOOT
- Dynamic
  - Dyninst
  - PIN
  - Valgrind

6

## Efficient Path Profiling

7

### Profiling (recap)

- Program profiling counts occurrences of an event during a program's execution
  - Basic blocks
  - Control-flow edges
  - Acyclic path
- Application
  - Performance tuning
  - Profile-directed compilation
  - Test coverage

8


## Goal

- Goal of paper: To efficiently collect path profiles for a DAG (i.e., acyclic-path profiling)
- Why not use existing techniques (existing at the time that the paper was written)?

## State of the Art

- Edge profiling: 16% overhead
- Estimation of path profiles from edge profiles  
(Correctly estimated only 38% of paths in SPEC benchmarks)

Explain this figure



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

## Acyclic-Path Profiling

- Assume for now—all paths acyclic—no loops
- Subsumes
  - basic block/statement profiling
  - edge/branch profiling
- Better approximation of intra-procedural path frequencies
- Stronger coverage criterion for white-box testing

12

## Goals for Instrumentation

- Low time and space overhead
- Minimal number of probes
- Optimal placement of the probes
- Paths represented with a simple integer value
- Compact numbering of paths


13

## Algorithm Overview (i)

- Each potential path is represented as a state
- Upon entry all paths are possible
- Each branch taken narrows the set of possible final states
- State reached at the end of the procedure represents the path taken

- Example:

- P0: 1, 2, 4, 5, 7
- P1: 1, 2, 4, 6, 7
- P2: 1, 3, 4, 5, 7
- P3: 1, 3, 4, 6, 7




14

## Algorithm Overview (i)

- Each potential path is represented as a state
- Upon entry all paths are possible
- Each branch taken narrows the set of possible final states
- State reached at the end of the procedure represents the path taken

- Example:

- P0: 1, 2, 4, 5, 7
- P1: 1, 2, 4, 6, 7
- P2: 1, 3, 4, 5, 7
- P3: 1, 3, 4, 6, 7



15

## Algorithm Overview (i)

- Each potential path is represented as a state
- Upon entry all paths are possible
- Each branch taken narrows the set of possible final states
- State reached at the end of the procedure represents the path taken


- Example:

P0: 1, 2, 4, 5, 7

P1: 1, 2, 4, 6, 7

P2: 1, 3, 4, 5, 7

P3: 1, 3, 4, 6, 7



16

## Algorithm Overview (ii)

- Final “states” (i.e., paths) are represented by integers in  $[0, n-1]$  ( $n$  == number of paths)
- Instrumentation not at every branch
- Transitions computed by simple arithmetic operations (no tables)
- CFG transformed in acyclic CFGs (DAGs)


- Example:

P0: 1, 2, 4, 5, 7

P1: 1, 2, 4, 6, 7

P2: 1, 3, 4, 5, 7

P3: 1, 3, 4, 6, 7



17



## Algorithm Steps


1. Assign integer values to edges such that no two paths compute the same path sum
2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge
3. Select appropriate instrumentation
4. Derive the executed paths from the collected run-time profiles

18

## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$


```
for each vertex  $v$  in rev. top. order {  
  if  $v$  is a leaf vertex {  
    NumPaths( $v$ ) = 1;  
  } else {  
    NumPaths( $v$ ) = 0;  
    for each edge  $e = v \rightarrow w$  {  
      Val( $e$ ) = NumPaths( $v$ );  
      NumPaths( $v$ ) += NumPaths( $w$ );  
    }  
  }  
}
```



19

## Topological, Reverse Topological Order

- Create a depth-first spanning tree
- Find topological and reverse topological orders
- Are there other orders based on other spanning trees?




20

## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```
for each vertex v in rev. top. order {  
  if v is a leaf vertex {  
    NumPaths(v) = 1;  
  } else {  
    NumPaths(v) = 0;  
    for each edge e = v->w {  
      Val(e) = NumPaths(v);  
      NumPaths(v) += NumPaths(w);  
    }  
  }  
}
```



21


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



i	Val(i)
1	NumPaths(n)

22


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



i	Val(i)
1	NumPaths(n)

23


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



i	Val(i)
Ⓜ	NumPaths(n)

24


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



i	Val(i)
Ⓜ	NumPaths(n)

25


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



$i$	$Val(i)$
$n$	$NumPaths(n)$

26


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



$i$	$Val(i)$
$n$	$NumPaths(n)$

27


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



$i$	$Val(i)$
$n$	$NumPaths(n)$

28


## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}

```



$i$	$Val(i)$
$n$	$NumPaths(n)$


29

## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}
    
```



$i$	$Val(i)$
$n$	$NumPaths(n)$


30

## Algorithm (Step 1 of 4)

1. Assign to each edge  $e$  a value  $Val(e)$  such that the sum along a path is unique and  $[0, n-1]$

```

for each vertex v in rev. top. order {
  if v is a leaf vertex {
    NumPaths(v) = 1;
  } else {
    NumPaths(v) = 0;
    for each edge e = v->w {
      Val(e) = NumPaths(v);
      NumPaths(v) += NumPaths(w);
    }
  }
}
    
```




Not necessarily the best placement

32


## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.



34

## Begin Side Discussion of Probe Placement



35




## Probe placement

- Knuth published efficient algorithms for finding the minimum number of edge counters for edge profiling
- Algorithm
  - Compute spanning tree  $T$  of CFG; edges in the spanning tree are bidirectional
  - Chords of spanning tree are edges  $E$  in CFG minus edges in  $T$
  - Instrumenting only the chords is sufficient to deduce execution of remaining edges

36

## Example

Show several spanning trees for the graph, and determine where probes should be placed (complete on board)



37

## Optimal Placement of Probes

- Several spanning trees may be possible on a CFG
- A maximum spanning tree is a spanning tree of maximum weight on its edges (why do we want maximum spanning tree?)
- Weight is defined as the execution frequency of the edge (how can this be done?)


38

## Edge Execution Frequency

- Program can be profiled to gather edge execution frequency
- Edge execution frequency can be approximated statically—static approximation heuristic [Ball and Larus 94]
- Generally impractical for purposes of profiling paths (why?)

39


## End Side Discussion of Probe Placement



40

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.




41

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY




42

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY
- Compute a maximal spanning tree (find chords)




43

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY
- Compute a maximal spanning tree (find chords)




44

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY
- Compute a maximal spanning tree (find chords)




45

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY
- Compute a maximal spanning tree (find chords)
- Assign increments: start from Val(e) and "propagate" to chord [Ball and Larus 94]




46

## Algorithm (Step 2 of 4)

2. Use a spanning tree to select edges to instrument and compute the appropriate increment for each instrumented edge.

- Add edge EXIT -> ENTRY
- Compute a maximal spanning tree (find chords)
- Assign increments: start from Val(e) and "propagate" to chord [Ball and Larus 94]




47

### Algorithm (Step 3 of 4)

#### 3. Select appropriate instrumentation

- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )



48

### Algorithm (Step 3 of 4)

#### 3. Select appropriate instrumentation


- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )

49

## Algorithm (Step 3 of 4)

### 3. Select appropriate instrumentation

- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )




50

## Algorithm (Step 3 of 4)

### 3. Select appropriate instrumentation

- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )
- Optimize
  - Initializations (first chord on paths)
  - Path's counter increment (last chord on paths)



51



## Algorithm (Step 3 of 4)

### 3. Select appropriate instrumentation

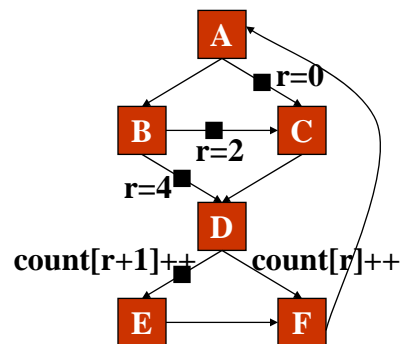
- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )
- Optimize
  - Initializations (first chord on paths)
  - Path's counter increment (last chord on paths)

52

## Algorithm (Step 3 of 4)

### 3. Select appropriate instrumentation

- Initialize path register ( $r=0$ )
- Update  $r$  in chords ( $r += \text{inc}$ )
- Increment path's counter at EXIT ( $\text{count}[r]++$ )
- Optimize
  - Initializations (first chord on paths)
  - Path's counter increment (last chord on paths)




53

### Algorithm (Step 4 of 4)

#### 4. Regenerating a path after collecting a profile

- Start at ENTRY
- Let  $r$  be the path value
- Select which edge to follow by finding the edge with the largest value  $Val(e) \leq r$
- Traverse edge  $e$  and  $r = r - Val(e)$




54

### Algorithm (Step 4 of 4)

#### 4. Regenerating a path after collecting a profile

- Start at ENTRY
- Let  $r$  be the path value
- Select which edge to follow by finding the edge with the largest value  $Val(e) \leq r$
- Traverse edge  $e$  and  $r = r - Val(e)$



**Generate path for 5**  
**Generate path for 3**

55


## Acyclic Paths

- All paths are intra-procedural (later extension to interprocedural)
- No cycles (to avoid infinite number of paths)
- Different kinds of loops and representations of loops (we'll see in what follows)

56

## Arbitrary Control Flow (loops)


- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register  
(`count[r]++`; `r=0`)



57

## Arbitrary Control Flow (loops)


- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register (count[r]++; r=0)
- This is not enough; with loops, 4 types of paths (v->w and x->y are back-edges)
  - ENTRY to EXIT
  - ENTRY to v (ending with execution of v->w)
  - w to x (after executing v->w and ending with the execution of x->y, v->w and x->y can be the same back-edge)
  - w to EXIT (after executing v->w)
- Need to distinguish them



58

## Arbitrary Control Flow (loops)


- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register (count[r]++; r=0)
- This is not enough; with loops, 4 types of paths (v->w and x->y are back-edges)
  - ENTRY to EXIT
  - ENTRY to v (ending with execution of v->w)
  - w to x (after executing v->w and ending with the execution of x->y, v->w and x->y can be the same back-edge)
  - w to EXIT (after executing v->w)
- Need to distinguish them



59

## Arbitrary Control Flow (loops)


- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register  
(`count[r]++`; `r=0`)
- This is not enough; with loops, 4 types of paths (`v->w` and `x->y` are back-edges)
  - ENTRY to EXIT
  - ENTRY to v (ending with execution of `v->w`)
  - w to x (after executing `v->w` and ending with the execution of `x->y`, `v->w` and `x->y` can be the same back-edge)
  - w to EXIT (after executing `v->w`)
- Need to distinguish them



60

## Arbitrary Control Flow (loops)


- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register  
(`count[r]++`; `r=0`)
- This is not enough; with loops, 4 types of paths (`v->w` and `x->y` are back-edges)
  - ENTRY to EXIT
  - ENTRY to v (ending with execution of `v->w`)
  - w to x (after executing `v->w` and ending with the execution of `x->y`, `v->w` and `x->y` can be the same back-edge)
  - w to EXIT (after executing `v->w`)
- Need to distinguish them



61

## Arbitrary Control Flow (loops)

- Loop implies the presence of a back-edge
- Back-edges instrumented to increment path counter and reinitialize path register (count[r]++; r=0)
- This is not enough; with loops, 4 types of paths (v->w and x->y are back-edges)
  - ENTRY to EXIT
  - ENTRY to v (ending with execution of v->w)
  - w to x (after executing v->w and ending with the execution of x->y, v->w and x->y can be the same back-edge)
  - w to EXIT (after executing v->w)
- Need to distinguish them



62

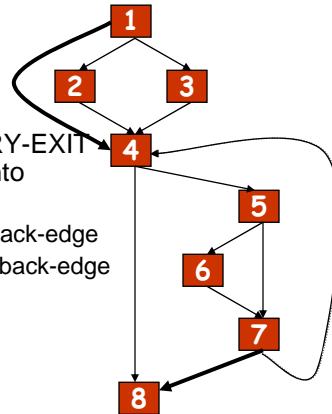
## Convert Arbitrary CFGs to DAGs

- Eliminate back-edges before computation of edge values and chord increments
- Remove a loop back-edge
- Add two edges
  - (1) ENTRY -> Target of back-edge
  - (2) Source of back-edge -> EXIT
- The dummy edges create extra paths ENTRY-EXIT that the value assignment algorithm takes into account
  - Edge (1) represents reinitializing along the back-edge
  - Edge (2) represents incrementing along the back-edge

63

## Convert Arbitrary CFGs to DAGs

- Eliminate back-edges before computation of edge values and chord increments
- Remove a loop back-edge
- Add two edges
  - (1) ENTRY -> Target of back-edge
  - (2) Source of back-edge -> EXIT
- The dummy edges create extra paths ENTRY-EXIT that the value assignment algorithm takes into account
  - Edge (1) represents reinitializing along the back-edge
  - Edge (2) represents incrementing along the back-edge



64

## Implementation

- Implemented in a tool called PP
- PP instruments SPARC binaries
- Built on top of EEL (binary instrumenter)
- Uses a register to store r
- Replaces array of counters with hash table if number of paths too large
- Plus some other optimizations

65

## Experimental Results (i)

- Used SPEC95 benchmark programs and test suites
- Edge profiling average overhead=16.1% (2.6%-52.8%)
- Path profiling average overhead=30.9% (5.5%-96.9%)
- When hashing is used performance is hurt
- Using no hashing, overhead is comparable or lower than edge profiling

66

## Algorithm Evolution

- Ball & Larus, "Optimally Profiling and Tracing Programs"
  - Focuses on edge and vertex profiling
  - Optimal placement of probes
- Ball, "Efficiently Counting Program Events with Support for On-line Queries"
  - Developed the technique for edge profiling with one register (instead of a counter for each edge)

69