#### Improving Software Reliability and Security via Symbolic Execution

#### **Cristian** Cadar

Department of Computing Imperial College London

Joint work with Paul Marinescu, Peter Collingbourne, Paul Kelly, JaeSeung Song, Peter Pietzuch, Hristina Palikareva (Imperial)

Dawson Engler, Daniel Dunbar, Junfeng Yang, Peter Pawlowski, Can Sar, Paul Twohey, Vijay Ganesh, David Dill, Peter Boonstoppel (Stanford)

Imperial College London

UPMARC Summer School Upsalla, 10-11 June 2013

### Writing Correct Software is Hard!

- Software complexity
  - Massive amounts of code
  - Tricky control flow
  - Complex dependencies
  - Intensive interaction w/ environment
- Code has to anticipate all possible interactions
  - Including malicious ones!

## Dynamic Symbolic Execution

Provides an automated way to reason about program behavior and the interaction with users and environment

- Dynamic symbolic execution can *automatically explore multiple paths* through a program
  - Using a constraint solver to determine the feasibility of each explored path
- Before each dangerous operation, can check if there are *any* values that can cause an error
- For each path, can usually generate a *concrete input triggering the path*

## Dynamic Symbolic Execution

- Toy example
- Scalability challenges
  - Path explosion, constraint solving challenges
- Generic bug finding
  - User-level utilities, kernel code, drivers, computer vision code, etc.
  - Attack generation against file systems and network servers
- Symbolic race detection for GPGPU code
- Semantic errors via crosschecking
  - Server interoperability, SIMD optimizations, GPU optimizations
- Patch testing
  - Testing six years of patches

#### Toy Example



#### All-Value Checks

## Implicit checks before each dangerous operation

- Pointer dereferences
- Array indexing
- Division/modulo operations
- Assert statements

#### All-value checks!

• Errors are found if **any** buggy values exist on that path!



#### All-Value Checks

## Implicit checks before each dangerous operation

- Pointer dereferences
- Array indexing
- Division/modulo operations
- Assert statements

#### All-value checks!

• Errors are found if **any** buggy values exist on that path!



## Mixed Concrete/Symbolic Execution

All operations that do not depend on the symbolic inputs are (essentially) executed as in the original code

#### Advantages:

- Ability to interact with the outside environment
  - E.g., system calls, uninstrumented libraries
- Can partly deal with limitations of constraint solvers
  - E.g., unsupported theories
- Only relevant code executed symbolically
  - Without the need to extract it explicitly

#### Three tools: EGT, EXE, KLEE



#### Scalability Challenges

## Path exploration challenges

Constraint solving challenges

## Path Exploration Challenges

Naïve exploration can easily get "stuck"

- Employing search heuristics
- Dynamically eliminating redundant paths
- Statically merging paths
- Using existing regression test suites to prioritize execution
- etc.

#### Search Heuristics

- Coverage-optimized search
  - Select path closest to an uncovered instruction
  - Favor paths that recently hit new code
- Best-first search
- Random path search
- etc.

#### Random Path Selection

- Maintain a binary tree of active paths
- Subtrees have equal prob. of being selected, irresp. of size
- NOT random state selection
- Favors paths high in the tree – fewer constraints
- Avoid starvation
  - e.g. symbolic loop



#### Which Search Heuristic?

Our latest tool KLEE uses multiple heuristics in a round-robin fashion, to protect against individual heuristics getting stuck in a local maximum.

#### Eliminating Redundant Paths

- If two paths reach the same program point with the same constraint sets, we can prune one of them
- We can discard from the constraint sets of each path those constraints involving memory which is never read again



#### Many Redundant Paths



#### Lots of Redundant Paths



#### **Redundant Path Elimination**



#### Statically Merging Paths





#### Statically Merging Paths

- Default: **2**<sup>N</sup> paths
- Phi-node folding: 1 path

#### **morph** computer vision algorithm: $2^{256} \rightarrow 1$

Path merging

 $\equiv$ 

Outsourcing problem to constraint solver

(which are often optimized for conjunctions of constraints)

## Using Existing Regression Suites

• Most applications come with a manually-written regression test suite

<pre>\$ cd lighttpd-1.4.29</pre>		
\$ make check		
•••		
./cachable.t	ok	
./core-404-handler.t	ok	
./core-condition.t	ok	
./core-keepalive.t	ok	
./core-request.t	ok	
./core-response.t	ok	
./core-var-include.t	ok	
./core.t	ok	
./lowercase.t	ok	
./mod-access.t	ok	

#### [ICSE'11]

#### **Regression Suites**

#### PROS

- Designed to execute interesting program paths
- Often achieve good coverage of different program features

#### CONS

- Execute each path with a single set of inputs
- Often exercise the general case of a program feature, missing corner cases

#### ZESTI: SymEx+Regression Suites

- 1. Use the paths executed by the regression suite to bootstrap the exploration process (to benefit from the coverage of the manual test suite and find additional errors on those paths)
- 2. Incrementally explore paths around the dangerous operations on these paths, in increasing distance from the dangerous operations (to test all possible corner cases of the program features exercised by the test suite)

[ICSE'11]

#### Multipath Analysis



#### Scalability Challenges

# Path exploration challenges

Constraint solving challenges

## **Constraint Solving Challenges**

#### **1. Accuracy:** need bit-level modeling of memory:

- Systems code often observes the same bytes in different ways: e.g., using pointer casting to treat an array of chars as a network packet, inode, etc.
- Bugs in systems code are often triggered by corner cases related to pointer/integer casting and arithmetic overflows

**2. Performance:** real programs generate many expensive constraints

## STP Constraint Solver [Ganesh, Dill]

- Modern constraint solver, based on *eager* translation to SAT (uses MiniSAT)
- Developed at Stanford by Ganesh and Dill, initially targeted to (and driven by) EXE
- Two data types: **bitvectors (BVs)** and **arrays of BVs**
- We model each memory block as an array of 8-bit BVs
- We can translate all C expressions into STP constraints with bit-level accuracy
  - Main exception: floating-point

#### **Constraint Solving: Accuracy**

- Mirror the (lack of) type system in C
  - Model each memory block as an array of 8-bit BVs
  - Bind types to expressions, not bits

# char buf[N]; // symbolic struct pkt1 { char x, y, v, w; int z; } \*pa = (struct pkt1\*) buf; struct pkt2 { unsigned i, j; } \*pb = (struct pkt2\*) buf; if (not of the optimized of the optimize

if (pa[2].v < 0) { assert(pb[2].i >= 1<<23); }



SBVLT (buf [18], 0x00)

BVGE (buf [19] @buf [18] @buf [17] @buf [16], 0x00800000)

#### **Constraint Solving: Performance**

- Inherently expensive (NP-complete)
- Invoked at every branch

• Key insight: exploit the characteristics of constraints generated by symex

#### Some Constraint Solving Statistics [after optimizations]

Application	Instrs/s	Queries/s	Solver %
[	695	7.9	97.8
base64	20,520	42.2	97.0
chmod	5,360	12.6	97.2
comm	222,113	305.0	88.4
csplit	19,132	63.5	98.3
dircolors	1,019,795	4,251.7	98.6
echo	52	4.5	98.8
env	13,246	26.3	97.2
factor	12,119	22.6	99.7
join	1,033,022	3,401.2	98.1
ln	2,986	24.5	97.0
mkdir	3,895	7.2	96.6
Avg:	196,078	675.5	97.1

1h runs using KLEE with DFS and no caching

UNIX utilites (and many other benchmarks)

- Large number of queries
- Most queries <0.1s
- Most time spent in the solver (before and after optimizations!)

[CAV'13]

## **Constraint Solving Optimizations**

Implemented at several different levels:

- SAT solvers
- SMT solvers
- Symbolic execution tools

#### Reasoning about Arrays in STP

• Many programs generate large constraints involving arrays with symbolic indexes

• STP handles this via array-based refinement

#### Reasoning about Arrays in STP

STP's conversion of array terms to SAT is expensive

$$(a[i_1] = e_1) \land (a[i_2] = e_2) \land (a[i_3] = e_3) \land (i_1 + i_2 + i_3 = 6)$$

$$(v_1 = e_1) \land (v_2 = e_2) \land (v_3 = e_3) \land (i_1 + i_2 + i_3 = 6)$$
  
 $(i_1 = i_2 \implies v_1 = v_2) \land (i_1 = i_3 \implies v_1 = v_3) \land (i_2 = i_3 \implies v_2 = v_3)$ 

Expands each formula by n·(n-1)/2 terms, where n is the number of syntactically distinct indexes

#### Array-based Refinement in STP

STP's conversion of array terms to SAT is expensive  $(a[i_1] = e_1) \land (a[i_2] = e_2) \land (a[i_3] = e_3) \land (i_1+i_2+i_3=6)$   $(v_1 = e_1) \land (v_2 = e_2) \land (v_3 = e_3) \land (i_1+i_2+i_3=6)$  $(i_1 = i_2 \implies v_1 = v_2) \land (i_1 = i_3 \implies v_1 = v_3) \land (i_2 = i_3 \implies v_2 = v_3)$ 



#### Array-based Refinement in STP

STP's conversion of array terms to SAT is expensive  $(a[i_1] = e_1) \wedge (a[i_2] = e_2) \wedge (a[i_3] = e_3) \wedge (i_1 + i_2 + i_3 = 6)$  $(v_1 = e_1) \wedge (v_2 = e_2) \wedge (v_3 = e_3) \wedge (i_1 + i_2 + i_3 = 6)$  $(i_1 = i_2 \longrightarrow v_1 = v_2) \land (i_1 = i_3 \longrightarrow v_1 = v_3) \land (i_2 = i_3 \longrightarrow v_2 = v_3)$  $i_{1} = 1$   $i_{2} = 2$   $i_{3} = 3$   $v_{1} = e_{1} = 1$   $v_{2} = e_{2} = 2$   $v_{2} = e_{2} = 3$  $(a[1] = 1) \land (a[2] = 2) \land$  $(a[3] = 3) \land (1+2+3 = 6)$
### Array-based Refinement in STP

STP's conversion of array terms to SAT is expensive  $(a[i_1] = e_1) \wedge (a[i_2] = e_2) \wedge (a[i_3] = e_3) \wedge (i_1 + i_2 + i_3 = 6)$  $(v_1 = e_1) \land (v_2 = e_2) \land (v_3 = e_3) \land (i_1 + i_2 + i_3 = 6)$  $\frac{(i_1 = i_2 \longrightarrow v_1 = v_2) \land (i_1 = i_3 \longrightarrow v_1 = v_3) \land (i_2 = i_3 \longrightarrow v_2 = v_3)}{(i_1 = i_3 \longrightarrow v_1 = v_3) \land (i_2 = i_3 \longrightarrow v_2 = v_3)}$  $\begin{bmatrix}
 i_1 = 2 \\
 i_2 = 2 \\
 i_3 = 2 \\
 v_1 = e_1 = 1 \\
 v_2 = e_2 = 2 \\
 v_3 = e_3 = 2
 \end{bmatrix}$  $(a[2] = 1) \land (a[2] = 2) \land$  $(a[2] = 3) \land (2+2+2 = 6)$ 

### Array-based Refinement in STP

STP's conversion of array terms to SAT is expensive  $(a[i_1] = e_1) \wedge (a[i_2] = e_2) \wedge (a[i_3] = e_3) \wedge (i_1 + i_2 + i_3 = 6)$  $(v_1 = e_1) \land (v_2 = e_2) \land (v_3 = e_3) \land (i_1 + i_2 + i_3 = 6)$  $(i_1 = i_2 \implies v_1 = v_2) \land (i_1 = i_3 \implies v_1 = v_3) \land (i_2 = i_3 \implies v_2 = v_3)$  $\begin{array}{c}
\overline{i_1 = 2} \\
i_2 = 2 \\
i_3 = 2 \\
v_1 = e_1 = 1 \\
v_2 = e_2 = 2 \\
\overline{v_2} = e_3 = 2
\end{array}$  $(a[2] = 1) \land (a[2] = 2) \land$  $(a[2] = 3) \land (2+2+2 = 6)$ 

#### Evaluation

Solver	Total time (min)	Timeouts
STP (baseline)	56	36
STP (array-based refinement)	10	1



8495 test cases from our

symbolic execution benchmarks

• Timeout set at 60s (which are added as penalty), underestimates performance differences

Higher-Level Constraint Solving Optimizations

- Two simple and effective optimizations
  - Eliminating irrelevant constraints
  - Caching solutions

## Eliminating Irrelevant Constraints

• In practice, each branch usually depends on a small number of variables



## **Caching Solutions**

• Static set of branches: lots of similar constraint sets



[OSDI'08] 42

### Speedup



### More on Caching: Instrs/Sec

Application	No caching	Caching	Speedup
[	3,914	695	0.17
base64	18,840	20,520	1.08
chmod	12,060	5,360	0.44
comm	73,064	222,113	3.03
csplit	10,682	19,132	1.79
dircolors	8,090	1,019,795	126.05
echo	227	52	0.22
env	21,995	13,246	0.60
factor	1,897	12,119	6.38
join	12,649	1,033,022	81.66
ln	13,420	2,986	0.22
mkdir	25,331	3,895	0.15
Avg:	16,847	196,078	11.63x

 Instrs/sec on ~1h runs, using DFS, w/ and w/o caching

> Need for better, more adaptive caching algorithms!

44

#### Portfolio of SMT Solvers



# EGT, EXE, KLEE

Successfully used our tools to:

- Automatically generate high-coverage test suites
- Find bugs and security vulnerabilities in complex software
- Perform bounded verification

#### Bug Finding with EGT, EXE, KLEE: Focus on Systems and Security Critical Code

	Applications
UNIX utilities	Coreutils, Busybox, Minix (over 450 apps)
UNIX file systems	ext2, ext3, JFS
Network servers	Bonjour, Avahi, udhcpd, lighttpd
Library code	libdwarf, libelf, PCRE, uClibc, Pintos
Packet filters	FreeBSD BPF, Linux BPF
MINIX device drivers	pci, lance, sb16
Kernel code	HiStar kernel
Computer vision code	OpenCV (filter, remap, resize, etc.)
OpenCL code	Parboil, Bullet, OP2

Most bugs fixed promptly

#### Coreutils Commands of Death

md5sum -c t1.txt	pr -e t2.txt
mkdir -Z a b	tac -r t3.txt t3.txt
mkfifo -Z a b	paste -d\\abcdefghijklmnopqrstuvwxyz
mknod -Z a b p	ptx -F\\abcdefghijklmnopqrstuvwxyz
seq -f %0 1	ptx x t4.txt
printf %d `	cut -c3-5,8000000output-d: file
t1.txt: \t t2.txt: \b\k	\tMD5(

[OSDI 2008, ICSE 2012]

### Attack Generation: File Systems

Some modern operating systems allow untrusted users to mount regular files as disk images!



## Attack Generation – File Systems

- Mount code is executed by the kernel!
- Attackers may create malicious disk images to attack a system

#### Attack Generation – File Systems



[Oakland 2006]

# Disk of death (JFS, Linux 2.6.10)

Offset	Hex Values							
00000	0000	0000	0000	0000	0000	0000	0000	0000
• • •	• • •							
08000	464A	3135	0000	0000	0000	0000	0000	0000
08010	1000	0000	0000	0000	0000	0000	0000	0000
08020	0000	0000	0100	0000	0000	0000	0000	0000
08030	E004	000F	0000	0000	0002	0000	0000	0000
08040	0000	0000	0000	• • •				

- 64<sup>th</sup> sector of a 64K disk image
- Mount it and PANIC your kernel

#### Attack Generation: Network Servers



### Bonjour: Packet of Death

Offset	Hex Values							
0000	0000	0000	0000	0000	0000	0000	0000	0000
0010	003E	0000	4000	FF11	1BB2	7F00	0001	E000
0020	OOFB	0000	14E9	002A	0000	0000	0000	0001
0030	0000	0000	0000	055F	6461	6170	045F	7463
0040	7005	6C6F	6361	6 <i>C</i> 00	000 <i>C</i>	0001		

- Causes Bonjour to abort, potential DoS attack
- Confirmed by Apple, security update released

Semantic Errors via Crosschecking (Equivalence Checking)

Lots of available opportunities as code is:Optimized frequentlyRefactored frequentlyDifferent implementations of the same interface



We can find any mismatches in their behavior by:

- 1. Using symbolic execution to explore multiple paths
- 2. Comparing the path constraints or input/output pairs across implementations

# Crosschecking: Advantages

- Can find semantic errors without the need for specifications!
- Constraint solving queries can be solved faster
- Can support constraint types not (efficiently) handled by the underlying solver, e.g., floating-point

Many crosschecking queries can be *syntactically* proven to be equivalent

### Crosschecking: Advantages



Many crosschecking queries can be *syntactically* proven to be equivalent

#### ZeroConf Protocol

- Enables devices to automatically configure themselves and their services and be discovered without manual intervention
- Two popular implementations: **Avahi** (opensource), and **Bonjour** (open-sourced by Apple)



### Server Interoperability Bonjour vs. Avahi

Offset	Hex Values							
0000	0000	0000	0000	0000	0000	0000	0000	0000
0010	003E	0000	4000	FF11	1BB2	7F00	0001	E000
0020	00FB	0000	14E9	002A	0000	0000	0002	0001
0030	0000	0000	0000	055F	6461	6170	045F	7463
0040	7005	6C6F	6361	6 <i>C</i> 00	000 <i>C</i>	0001		

- mDNS specification (§18.11): *"Multicast DNS messages received with non-zero Response Codes MUST be silently ignored."*
- Avahi ignores this packet, Bonjour does NOT

### New Platforms, New Code

- Recent years have seen the emergence of new computing platforms which provide many opportunities for optimizations
- Code is often adapted manually to benefit from these platforms

*Error-prone*, as any manual process

# **SIMD** Optimizations

Most processors offer support for SIMD instructions

- Can operate on multiple data concurrently
- Many algorithms can make use of them (e.g., computer vision algorithms)



[EuroSys 2011]

# OpenCV

Popular computer vision library from Intel and Willow Garage



[Corner detection algorithm]

Computer vision algorithms were optimized to make use of SIMD



# **OpenCV** Results

- Crosschecked 51 SIMD-optimized versions against their reference scalar implementations
  - Proved the bounded equivalence of 41
  - Found mismatches in 10
- Most mismatches due to tricky FP-related issues:
  - Precision
  - Rounding
  - Associativity
  - Distributivity
  - NaN values

### **OpenCV** Results

Surprising find: min/max not commutative nor associative!

min(a,b) = a < b ? a : b

a < b (ordered) → always returns false if one of the operands is NaN

min(NaN, 5) = 5 min(5, NaN) = NaN

min(min(5, NaN), 100) = min(NaN, 100) = 100 min(5, min(NaN, 100)) = min(5, 100) = 5

### **GPU** Optimizations



## General Purpose GPU Computing



# General Purpose GPUs (GPGPUs)



General-Purpose Graphics Processing Units (GPGPUs) are a programmable platform for highly parallel computation

New programming model:

- Large number of threads
- Hierarchical execution and memory model

# OpenCL

- Open Computing Language (OpenCL): an open standard for parallel computation

   Targets both CPUs and GPGPUs
- OpenCL C language is a dialect of C99
- An OpenCL C program consists of one or more OpenCL C kernels which are executed on a device (such as a GPU)
- OpenCL C kernels may be invoked in parallel by the host using multiple work-items

### OpenCL Example



work-items

### **Race Conditions**

```
kernel void arr sqrt(global float *a) {
   size t i = get global id(0),
          N = get local size(0);
   float r0 = i > 0 ? a[i-1]:0;
   float r1 = a[i];
   float r2 = i < N ? a[i+1]:0;
   a[i] = (r0+r1+r2)/3;
```

```
R/W race: work-item i writes a[i],
work-item i+1 reads a[i]
```

#### Barriers

```
kernel void arr sqrt(global float *a) {
   size t i = get global id(0),
          N = get local size(0);
   float r0 = i > 0 ? a[i-1]:0;
   float r1 = a[i];
   float r2 = i < N ? a[i+1]:0;</pre>
   barrier(CLK GLOBAL MEM FENCE);
   a[i] = (r0+r1+r2)/3;
```

barrier() blocks until all work-items
(in the same work-group) reached the call

### Race Checking for GPGPUs

#### MAR(a[0]): work item 0

	Wid	R	W
kernel void arr_sqrt(global float *a) {	-	-	-
<pre>size_t i = get_global_id(0),</pre>	-	-	-
N = get_local_size(0);	-	-	-
float $r0 = i > 0$ ? $a[i-1]:0;$	-	-	-
float $r1 = a[i];$	0	X	-
float $r^2 = i < N ? a[i+1]:0;$	0	Χ	-
a[i] = (r0+r1+r2)/3;	0	X	X
}			

- Wid: Work-item that accessed byte
- R: whether byte was read
- W: whether byte was written
## Race Checking for GPGPUs

#### MAR(a[0]): work item 1

	Wid	R	W	
loat *a) {	0	Х	Χ	
),	0	X	Χ	
0);	0	X	Χ	
0;	0	X	X	R/W
0;				

```
_kernel void arr_sqrt(global float *a) {
    size_t i = get_global_id(0),
        N = get_local_size(0);
    float r0 = i > 0 ? a[i-1]:0;
    float r1 = a[i];
    float r2 = i < N ? a[i+1]:0;
    a[i] = (r0+r1+r2)/3;
}</pre>
```

- Wid: Work-item that accessed byte
- R: whether byte was read
- W: whether byte was written

### Race Checking for GPGPUs

#### MAR(a[0]): work item 0

	Wid	R	W
kernel void arr_sqrt(global float *a) {	-	-	-
<pre>size_t i = get_global_id(0),</pre>	-	-	-
N = get_local_size(0);	-	-	-
float $r0 = i > 0$ ? $a[i-1]:0;$	-	-	-
float $r1 = a[i];$	0	X	-
float $r2 = i < N ? a[i+1]:0;$	0	X	-
<pre>barrier(CLK_GLOBAL_MEM_FENCE);</pre>			
a[i] = (r0+r1+r2)/3;			
}			

We model barrier() by resetting the MAR before continuing execution any of the work-items past the barrier

### Race Checking for GPGPUs

#### MAR(a[0]): work item 1

	Wid	R	V
<pre>kernel void arr_sqrt(global float *a) {</pre>	0	X	-
<pre>size_t i = get_global_id(0),</pre>	0	X	-
N = get_local_size(0);	0	X	-
float $r0 = i > 0$ ? $a[i-1]:0;$	0	X	-
float $r1 = a[i];$	0	X	-
float $r2 = i < N ? a[i+1]:0;$	0	X	-
<pre>barrier(CLK_GLOBAL_MEM_FENCE);</pre>			
a[i] = (r0+r1+r2)/3;			

We model barrier() by resetting the MAR before continuing execution any of the work-items past the barrier

#### Symbolic Races



Write-after-write race if **i=j** satisfiable

# GPGPU (OpenCL) Optimizations

- Parboil:
  - GPU benchmark suite, originally written in CUDA
- OP2
  - Library for applications on unstructured grids
- Bullet open-source physics library
  - Popular library used movie studios and professional game developers
  - Analyzed soft body engine

**Bullet library** 



OpenCL Benchmarks: Bugs and Mismatches

Several bugs and mismatches:

- 2 mismatches between C and OpenCL code
  - Incorrect FP associativity and distributivity assumptions (CP in Parboil)
- 3 memory errors
  - Buffer overflows (MRI-Q&MRI-FHD in Parboil)
  - Use-after-free: incorrect synchronization between host and kernel code (MRI-Q in Parboil)
  - Uninitialized memory (MRI-FHD in Parboil)
- 1 race condition
  - Missing synchronization barrier (OP2)
- 1 compiler bug
  - NVidia compiler bug (incorrect optimization)

## Integrating Crosschecking into Development Process

Semantic mismatches not always errors

- Underspecified behavior
- Two (anecdotal) insights:
- 1. Provide developers the **ability to add "assumptions"** eg:
  - Floating-point associativity holds:
    - A+(B+C) = (A+B)+C
  - Disregard the difference between  $0_{-}$  and  $0_{+}$ :
    - A+0 = A
- 2. All things being equal, developers **prefer to keep the behavior of the reference implementation** 
  - Particularly if we can provide some guarantees
    - bounded equivalence

#### **KATCH: High-Coverage Symbolic Patch Testing**

test<sub>₄</sub>

 $rest_4$ 

test<sub>4</sub>

test<sub>4</sub>

bug









# Greedy Exploration Step

```
+ char escinput = escape(input);
```

```
+ write(file, &escinput, 1);
```

```
lighttpd r2660: patch
modifies log() to escape
sensitive characters
```

close(file);

+ }

# Greedy Exploration Step

```
void log(char input) {
   int file = open("access.log"...);
   if (input >= '_' && input <= '~') {
      // printable characters
      write(file, &input, 1);
 + } else {
      char escinput = escape(input);
 +
      write(file, &escinput, 1);
 + }
   close(file);
```

Available input: "t" (or any printable char)

 Greedy step: choose the symbolic branch point whose unexplored side is closest to the patch.
 Explore this side!

#### Informed Path Regeneration

void log(char input) { if (input >= '_' && input <= '~') {	Avail
 } else {	G
+	
}	1. B
}	sy
	th
if (0 == strcmp(request, "GET")	si
	2. E
for (char* p = request; *p; p++) log(*p);	si

Available input: "GET"

**Greedy step fails!** 

- Backtrack to the last symbolic branch that disallows this side to be executed
- 2. Explore the other side of that branch

req[2] ≠ 'T'

## Definition Switching

<pre>void log(char input) {     if (escape == ESCAPE_ALL) {</pre>	}	<b>Backtracking step fails!</b>	
<pre>void log(char input) {     if (escape == ESCAPE_ALL) {         Patch guarded by         an anata here also         ALL) {</pre>	+	concrete branch	
void log(char input) { if (access = ECCAPE ALL) ( Patch guarded by	If (escape == ESCAPE_ALL) {	aananata hranah	
void loo(obon input) (	if (access = ECAPE ALL) (	Patch guarded by	
	void loc(chan innut) (		
enum escape_t escape; Available test: ont = 'a	enum escape_t escape;	Available test: $opt = {a'}$	

```
opt = getopt_long(argc, argv, ...);
switch (opt) {
    case 'a': escape = ESCAPE_SPACE;
        break;
    case 'b': escape = ESCAPE_ALL;
```

log(...);

- 1. Find all reaching definitions for the variables involved and try to cover another one.
- 2. Favors definitions that can be statically shown to satisfy target, or unexecuted definitions

### Input Selection

Naïve solution: calculate the context-sensitive static distance between the path executed by an input and the patch code

if (x < 100) f(x); else if (x > 200) f(x+1); void f(int x) { if (x % 2 == 0) PATCH;



## Input Selection

Naïve solution: calculate the context-sensitive static distance between the path executed by an input and the patch code

if (x < 100) f(x); else if (x > 200) f(x+1); void f(int x) { if (x == 999) PATCH:



### Input Selection: Weakest Preconditions

For which basic block in the program, compute a necessary condition for reaching the target.

Prune CFG edges which make the target unreachable.

if (x < 100)
 f(x);
else
 if (x > 200)
 f(x+1);
void f(int x) {
 if (x == 999)
 PATCH;





Revision	ELOC	Covered ELOC	
		Regression	KATCH
2631	20	15 (75%)	20 (100%)

http://zzz.example.com/



https://zz.example.com/

91

## Lighttpd r2660

Revision	ELOC	Covered ELOC	
		Regression	KATCH
2660	33	9 (27%)	24 (72%)

- 165 if (str->ptr[i] >= '\_' && str->ptr[i] <= '~') {</pre>
- 166 /\* printable chars \*/
- 167 buffer\_append\_string\_len(dest,&str ->ptr[i],1);
- 168 } else switch (str->ptr[i]) {
- 169 case '"':
- 170 BUFFER\_APPEND\_STRING\_CONST(dest, "\\\"");
- 171 break;

Bug reported and fixed promptly by developers



Key evaluation criteria: **no cherry picking!** 

• choose all patches for an application over a contiguous time period

FindUtils suite (FU) find, xargs, locate	12,648 ELOC	125 patches written over ~26 months
DiffUtils suite (DU)	55,655 ELOC	175 patches written
s/diff, diff3, cmp	+ 280,000 in libs	over ~30 months
<b>BinUtils suite (BU)</b>	81,933 ELOC	181 patches written
ar, elfedit, nm, etc.	+ 800,000 in libs	over ~16 months

## Patch Coverage (basic block level)



Standard symbolic execution (30min/BB) only added +1.2% to FU

## Patch Coverage (current results – ongoing work)



Standard symbolic execution (30min/BB) only added +1.2% to FU

# Binutils: Coverage+Bugs (current results)



- Found 14 distinct crash bugs, all in BU
  - All unreachable by standard symbolic execution
- 12 bugs still present in latest version
  - Reported (some already fixed) by dev '
- 10 bugs found in the patch code itself or affected by patch code

#### KLEE – Demo

#### http://klee.llvm.org

- bad\_abs
- xcheck\_abs
- squares

#### **Dynamic Symbolic Execution**

- Automatically reasons about program behavior and the interaction with users and environment
- Can generate inputs exposing both generic and semantic bugs in complex software
  - Including file systems, library code, utility applications, network servers, device drivers, computer vision code

#### KLEE: Freely Available as Open-Source

#### http://klee.llvm.org

- Over 250 subscribers to the klee-dev mailing list
- Extended in many interesting ways by several research groups, in the areas of:
  - wireless sensor networks/distributed systems
  - schedule memoization in multithreaded code
  - automated debugging
  - exploit generation
  - client-behavior verification in online gaming
  - GPU testing and verification
  - etc.