

# An Introduction to Dynamic Symbolic Execution and the KLEE Infrastructure

**Cristian Cadar**

**Department of Computing  
Imperial College London**



**SOFTWARE RELIABILITY  
GROUP**

# Dynamic Symbolic Execution

---

- Dynamic symbolic execution is a technique for *automatically exploring paths* through a program
  - Determines the feasibility of each explored path using a *constraint solver*
  - Checks if there are *any* values that can cause an error on each explored path
  - For each path, can generate a *concrete input triggering the path*

# Dynamic Symbolic Execution

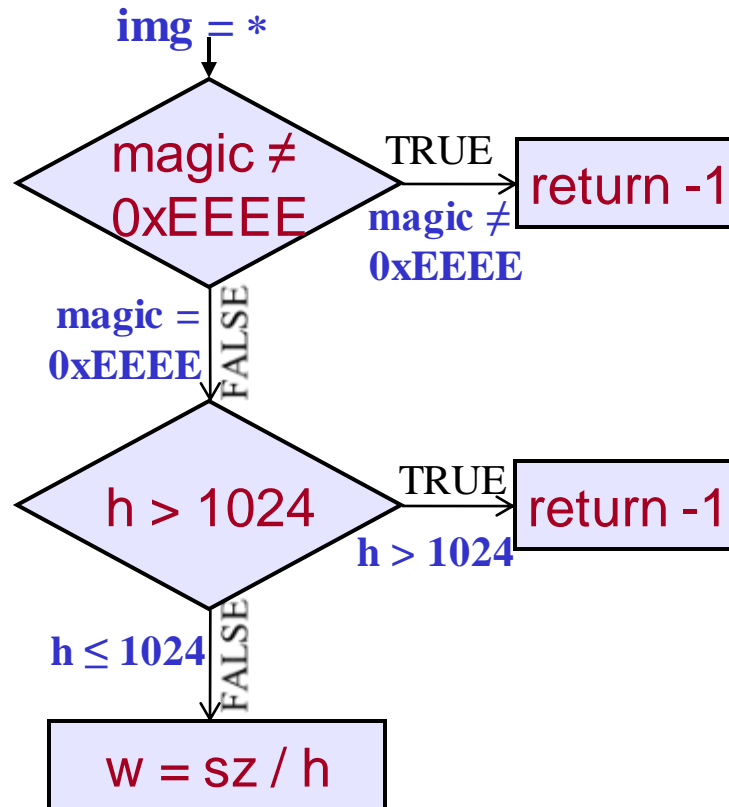
---

- Received significant interest in the last few years
- Many dynamic symbolic execution/concolic tools available as open-source:
  - **CREST, KLEE, SYMBOLIC JPF**, etc.
- Started to be adopted/tried out in the industry:
  - Microsoft (**SAGE, PEX**)
  - NASA (**SYMBOLIC JPF, KLEE**)
  - Fujitsu (**SYMBOLIC JPF, KLEE/KLOVER**)
  - IBM (**APOLLO**)
  - etc.

# Toy Example

```
struct image_t {  
    unsigned short magic;  
    unsigned short h, sz;  
    ...  
}
```


```
int main(int argc, char** argv) {  
    ...  
    image_t img = read_img(file);  
    if (img.magic != 0xEEEE)  
        return -1;  
    if (img.h > 1024)  
        return -1;  
    w = img.sz / img.h;  
    ...  
}
```



# Toy Example

```
struct image_t {  
    unsigned short magic;  
    unsigned short h, sz;  
    ...  
}
```

```
int main(int argc, char** argv) {  
    ...  
    image_t img = read_img(file);  
    if (img.magic != 0xEEEE)  
        return -1;  
    if (img.h > 1024)  
        return -1;  
    w = img.sz / img.h;  
    ...  
}
```



# 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!

```
int foo(unsigned k) {  
    int a[4] = {3, 1, 0, 4};  
    k = k % 4;  
    return a[a[k]];  
}
```



# 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!

```
int foo(unsigned k) {  
    int a[4] = {3, 1, 0, 4};  
    k = k % 4;  
    return a[a[k]];  
}
```



Buffer overflow!

# 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



# KLEE

---

- Symbolic execution tool started as a successor to EXE
- Based on the LLVM compiler, primarily targeting C code
- Open-sourced in June 2009, now available on GitHub
- Active user base with over 300 subscribers on the mailing list and over 35 contributors listed on GitHub

Webpage: **[klee.github.io](http://klee.github.io)**

Code: **<https://github.com/klee>**

# KLEE

---


- Extensible platform, used and extended by many groups in academia and industry, in the areas such as:
  - bug finding
  - high-coverage test input generation
  - exploit generation
  - automated debugging
  - wireless sensor networks/distributed systems
  - schedule memoization in multithreaded code
  - client-behavior verification in online gaming
  - GPU testing and verification, etc.

An incomplete list of publications and extensions available at:

**[klee.github.io/Publications.html](http://klee.github.io/Publications.html)**

# High Line Coverage

(Coreutils, non-lib, 1h/utility = 89 h)



[Cadaru, Dunbar, Engler OSDI 2008]

# Bug Finding with KLEE (incl. EGT/EXE):

## 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, etc.
Library code	libdwarf, libelf, PCRE, uClibc, etc.
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
```

```
mkdir -Z a b
```

```
mkfifo -Z a b
```

```
mknod -Z a b p
```

```
seq -f %0 1
```

```
printf %d `
```

```
pr -e t2.txt
```

```
tac -r t3.txt t3.txt
```

```
paste -d\ abcdefghijklmnopqrstuvwxyz
```

```
ptx -F\ abcdefghijklmnopqrstuvwxyz
```

```
ptx x t4.txt
```

```
cut -c3-5,8000000- --output-d: file
```

*t1.txt:* \t \tMD5 (

*t3.txt:* \n

*t2.txt:* \b\b\b\b\b\b\b\b\b

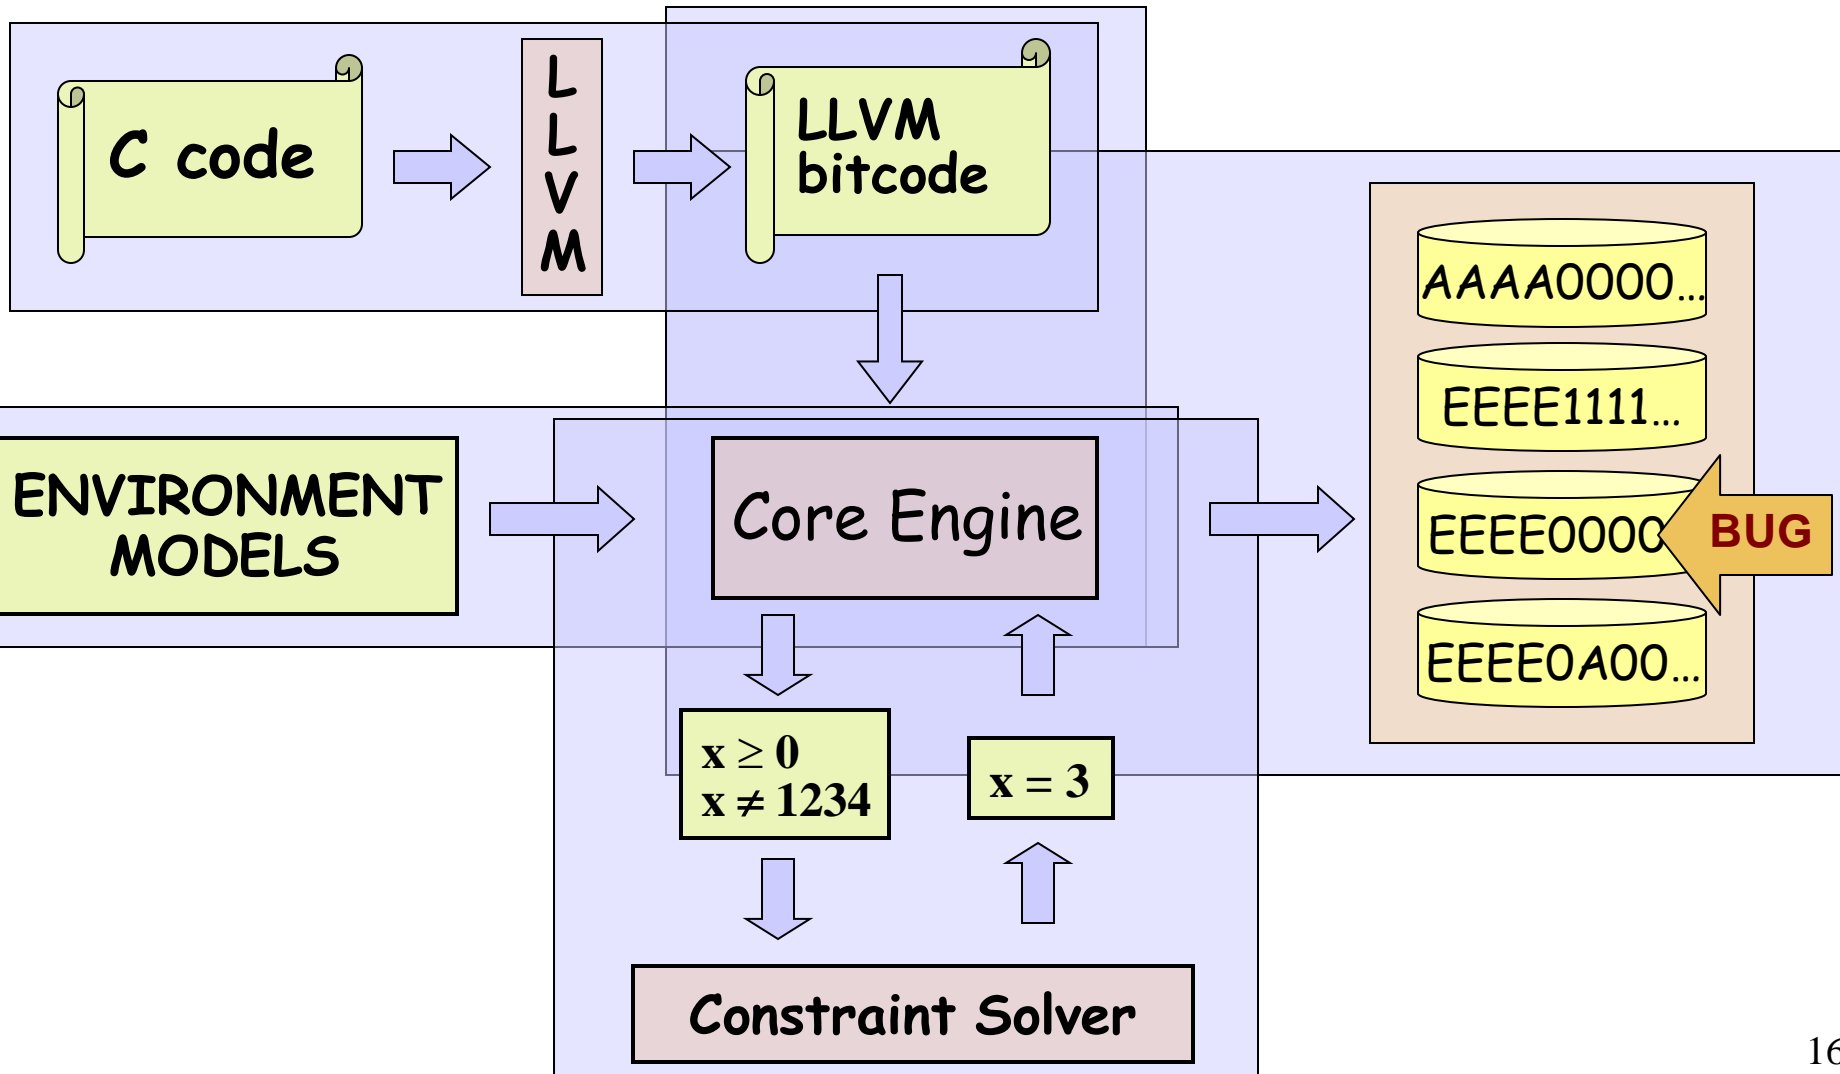
*t4.txt:* A

# Packet of Death (Bonjour)

Offset	Hex Values								
0000	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	0000	0001	
0030	0000	0000	0000	055F	6461	6170	045F	7463	
0040	7005	6C6F	6361	6C00	000C	0001			

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

# KLEE Architecture



# KLEE Demo: Toy Image Viewer

```
// #include directives
struct image_t {
    unsigned short magic;
    unsigned short h, sz; // height, size
    char pixels[1018];
};

int main(int argc, char** argv) {
    struct image_t img;
    int fd = open(argv[1], O_RDONLY);
    read(fd, &img, 1024);

    if (img.magic != 0xEEEE)
        return -1;
    if (img.h > 1024)
        return -1;
    unsigned short w = img.sz / img.h;

    return w;
}
```

```
$ clang -emit-llvm -c -g image_viewer.c
$ klee --posix-runtime -write-pcs
    image_viewer.bc --sym-files 1 1024 A
...
KLEE: output directory = klee-out-1
(klee-last)
...
KLEE: ERROR: ... divide by zero
...
KLEE: done: generated tests = 4
```



# KLEE Demo: Toy Image Viewer

---

```
$ cat klee-last/test000003.pc
...
array A-data[1024] : w32 -> w8 = symbolic
(query [
  ...
  (Eq 61166
    (ReadLSB w16 0 A-data))
  (Eq 0
    (ReadLSB w16 2 A-data))
  ...
])
```

# KLEE Demo: Toy Image Viewer

---

```
$ klee-replay --create-files-only klee-last/test000003.ktest  
[File A created]
```

```
$ xxd -g 1 -l 10 A  
0000000: ee ee 00 00 00 00 00 00 00 00 00 .....  
.....
```

```
$ gcc -o image_viewer image_viewer.c  
[image_viewer created]
```

```
$ ./image_viewer A  
Floating point exception
```

# KLEE Demo: All-Values Checks

---

```
int foo(unsigned k) {
    int a[4] = {3, 1, 0, 4};
    k = k % 4;
    return a[a[k]];
}

int main() {
    int k;
    klee_make_symbolic(&k, sizeof(k), "k");
    return foo(k);
}
```

```
$ clang -emit-llvm -c -g all-values.c
$ klee all-values.bc
...
KLEE: ERROR: /home/klee/all-values/all-values.c:4: memory error: out of bound pointer
...
KLEE: done: completed paths = 2
KLEE: done: generated tests = 2
```

# Running KLEE inside a Docker container

---

**Step 1:** Install Docker for Linux/MacOS/Windows

**Step 2:** `docker pull klee/klee`

**Step 3:** `docker run --rm -ti --ulimit='stack=-1:-1' klee/klee`

**<http://klee.github.io/docker/>**


# KLEE on the Web

---


You can try KLEE on the web now (world premiere!) at:

**<http://klee.doc.ic.ac.uk>**  
(work in progress)

# KLEE Architecture




# KLEE Architecture:



LLVM advantages:

- Mature framework, incorporated into commercial products by Apple, Google, Intel, etc.
- Elegant design patterns: analysis passes, visitors, etc.
- Single Static-Assignment (SSA) form with infinite registers (nice fit for symbolic execution)
- Lots of useful program analyses
- Well documented
- Several different front-ends, so KLEE could be extended to work with languages other than C

# KLEE Architecture:




## LLVM disadvantages

- Fast changing, not-backward compatible API!
  - KLEE is currently based on LLVM 3.4
- Compiling to LLVM bitcode is still not trivial, but it's getting better:
  - `make CC="clang -emit-llvm"`
  - LLVM Gold Plugin <http://llvm.org/docs/GoldPlugin.html>
  - Whole-Program LLVM <https://github.com/travitch/whole-program-llvm>



# KLEE Architecture:



KLEE runs LLVM, not C code!

```
#include <stdio.h>
int main() {
    int x;
    klee_make_symbolic(&x, sizeof(x), "x");

    if (x > 0)
        printf("x\n");
    else printf("x\n");


    return 0;
}
```

```
$ clang -emit-llvm -c -g code.c
$ klee code.bc
...
x
KLEE: done: total instructions = 6
KLEE: done: completed paths = 1
KLEE: done: generated tests = 1
```

# KLEE Architecture:

Core Engine

The core engine implements symbolic execution exploration.



# KLEE Architecture:


Interpreter

Core Engine

- Works as a mixed concrete/symbolic interpreter for LLVM bitcode

```
Instruction *i = ki->inst;
switch (i->getOpcode()) {
  case Instruction::Ret:
    ...
  case Instruction::Br:
    // if both sides feasible, fork
    ...
}
```

\$ ./program



\$ klee program.bc



# KLEE Architecture:

Core Engine

The core engine implements symbolic execution exploration.

Two main **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 in KLEE

Core Engine

Searchers

- Basic search heuristics such as BFS and DFS

```
klee --search=bfs program.bc
```

- Coverage-optimized search (**--search=nurs:md2u**)

- Select path closest to an uncovered instruction

- Random-state search (**--search=random-state**)

- Randomly select a pending state/path

- Random-path search (**--search=random-path**)

- Described next

- etc.

[Cadar, Ganesh, Pawlowski, Dill, Engler CCS'06]

[Cadar, Dunbar, Engler OSDI'08]


[Marinescu, Cadar ICSE'12], etc.

# Random Path Selection

Core Engine

Searchers

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





# Combining Search Heuristics

Core Engine

Searchers

KLEE can also use multiple heuristics in a round-robin fashion, to protect against individual heuristics getting stuck in a local maximum.


```
klee --search=nurs:md2u --search=dfs  
--search=random-path ...
```

# New Search Heuristics

Core Engine

Searchers


Tree of ESs



Easy to plug a new searcher by extending the Searcher class:

```
selectState() → ExecutionState  
update(addedStates, removedStates)
```

CFG



Statistics

- Solver time
- Instructions executed
- Memory consumption
- etc.

# Memory Modelling

Core Engine

Memory

**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

# Memory Modelling

Core Engine

Memory

- One data type: **arrays of bitvectors (BVs)**
- 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
- We can translate all C expressions into constraints in the theory of quantifier-free BV with arrays (QF\_ABV) with bit-level accuracy
  - Main exception: floating-point

# Accuracy: Example

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

```
buf: ARRAY BITVECTOR(32) OF BITVECTOR(8)
```

```
buf[18] <SIGNED 0x00
```

```
buf[19]@buf[18]@buf[17]@buf[16] ≥UNSIGNED 0x00800000
```

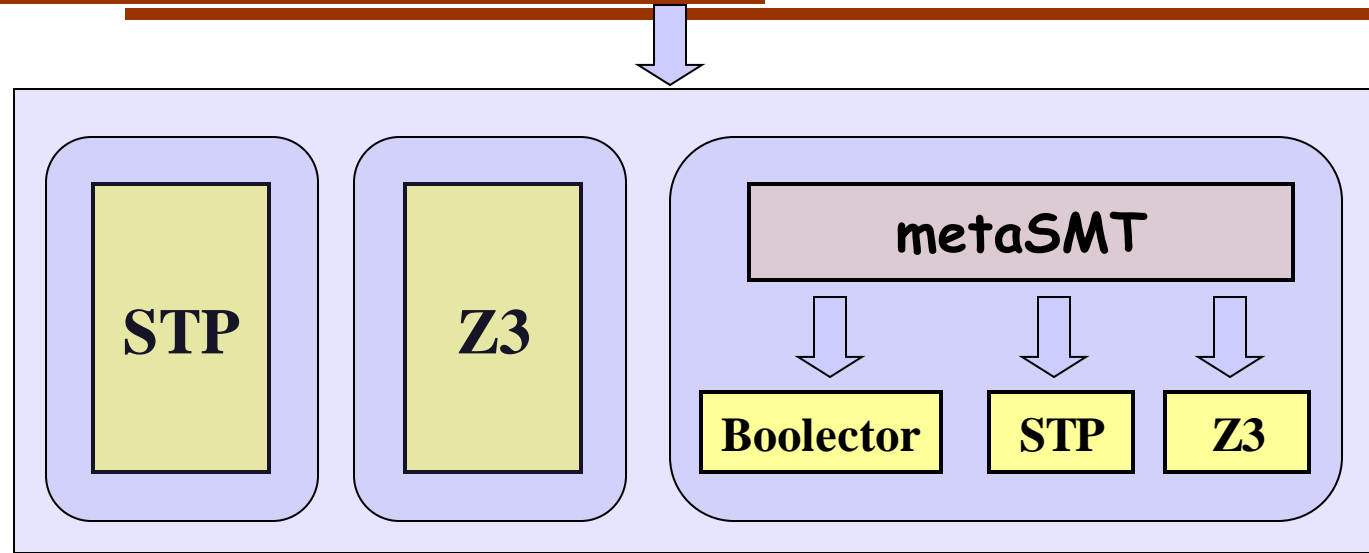
# KLEE Architecture



# SMT Solvers


(--solver-backend=stp, z3, ...)

Theory of closed  
quantifier-free  
formulas over  
bitvectors and  
arrays of  
bitvectors  
(QF\_ABV)



- **STP**: Developed at Stanford. Initially targeted to, and driven by, EXE. Main solver in KLEE.
- **Z3**: Developed at Microsoft Research, integrated both natively and as part of metaSMT.
- **Boolector**: Developed at Johannes Kepler University, integrated via metaSMT.

# metaSMT



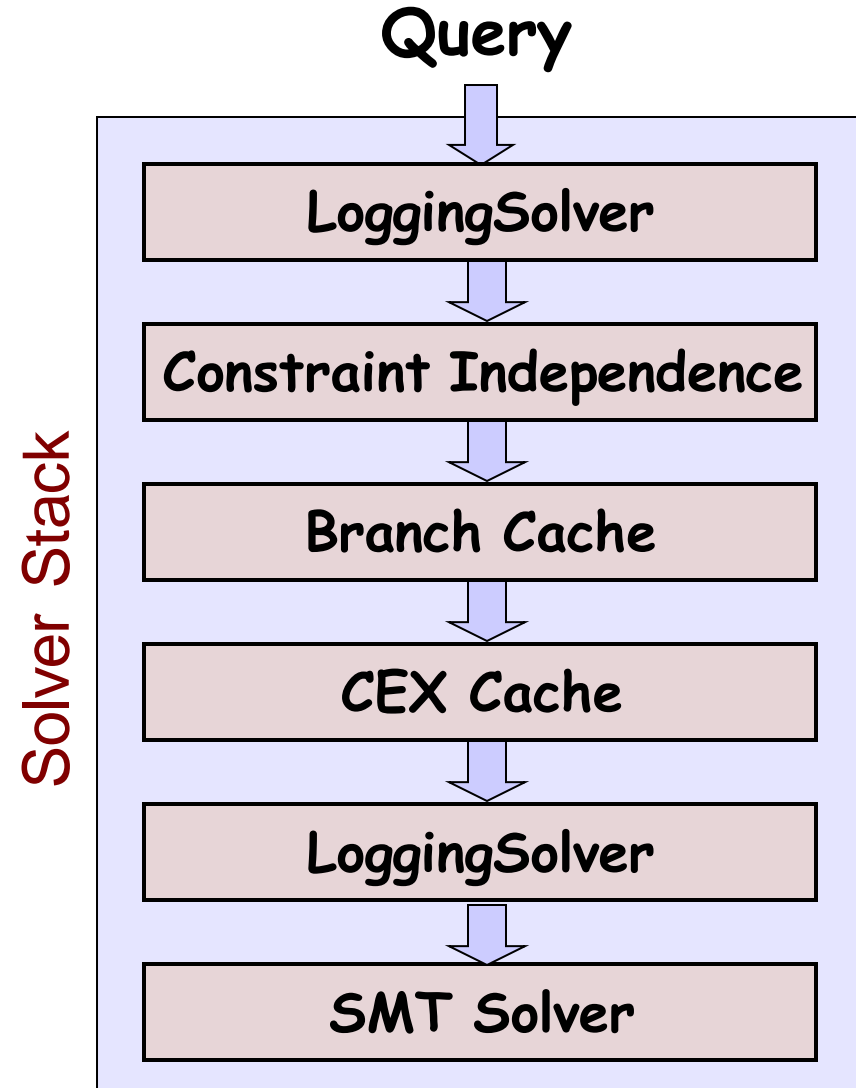
- **metaSMT** developed at University of Bremen provides a unified API for transparently using a number of SMT (and SAT) solvers
  - Avoids communication via text files, which would be too expensive
  - Small overhead: compile-time translation via metaprogramming



# KLEE Architecture:

Constraint Solver

- Several high-level optimizations specific to symex
  - CEX caching, elimination of irrelevant constraints, etc.
- Implemented as a stack of solver passes
- Caching → only some queries reach the solver
- Independent **Kleaver** tool that implements this solver stack



# Constraint Solving: Performance

---

- Inherently expensive
- 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
<b>Avg:</b>	<b>196,078</b>	<b>675.5</b>	<b>97.1</b>

1h runs using KLEE with  
DFS and no caching

UNIX utilities (and many  
other benchmarks)

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

# Higher-Level Constraint Solving Optimizations

---

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

# Eliminating Irrelevant Constraints

(`--use-independent-solver=true/false`)

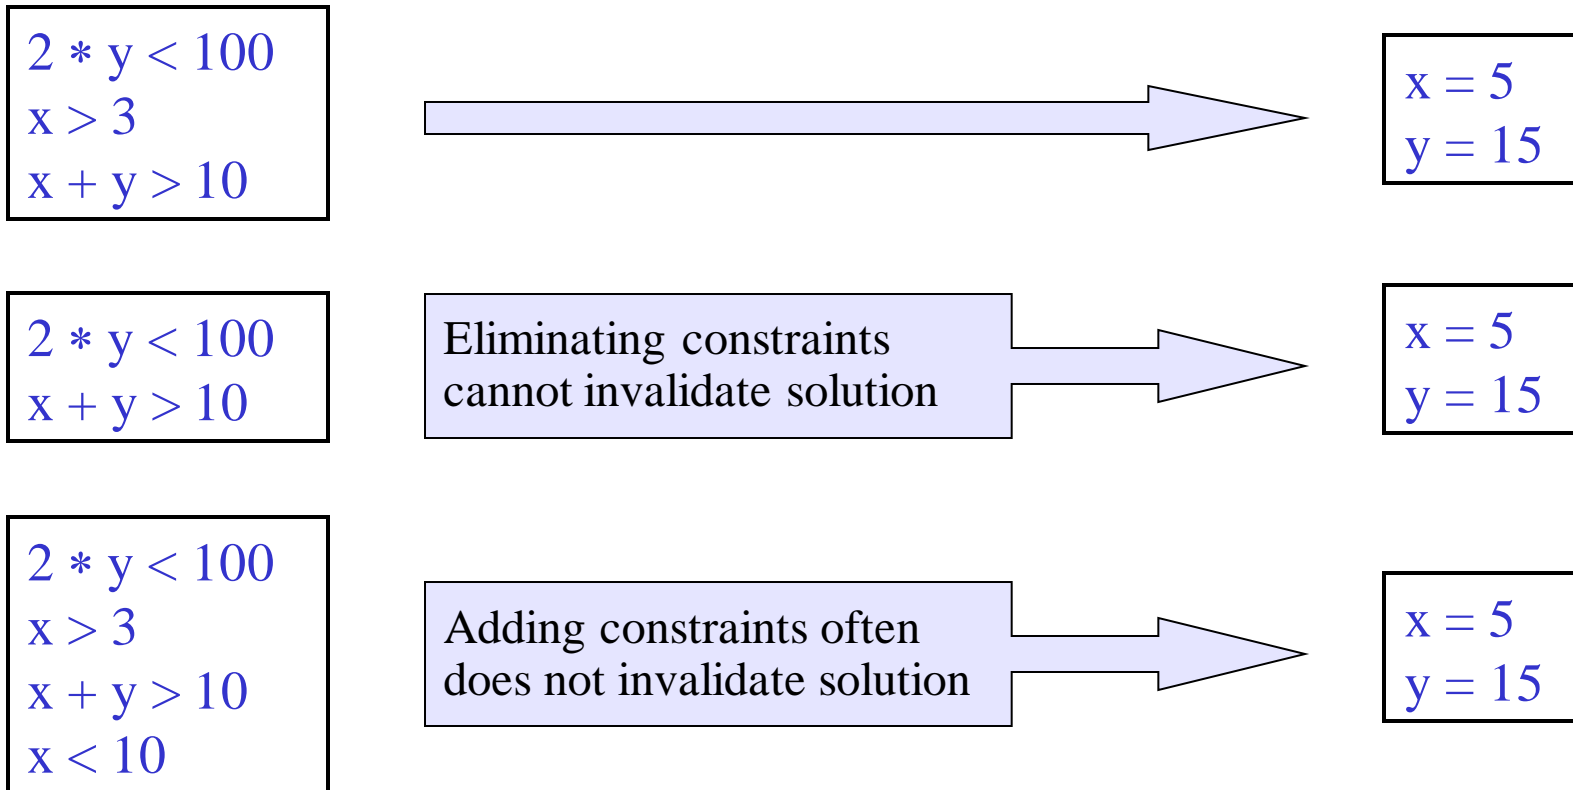
- In practice, each branch usually depends on a small number of variables

...	<del><math>w + z &gt; 100</math></del>
...	<del><math>2 * w - 1 &lt; 12345</math></del>
...	$x + y > 10$
...	<del><math>z \&amp; -z = z</math></del>
if (x < 10) {	$x < 10 ?$
...	
}	


# Caching Solutions

(`--use-cex-cache=true/false`)


- Static set of branches: lots of similar constraint sets



# Speedup



# KLEE Architecture





# KLEE Architecture:

## Environment Models

- Environment model: model for a piece of code for which source is not available
- In KLEE, the environment is mainly the OS system call API

# Environmental Modeling

*Models are plain C code, which KLEE interprets as any other code!*

```
// actual implementation: ~50 LOC
ssize_t read(int fd, void *buf, size_t count) {
    klee_file_t *f = get_file(fd);
    ...
    memcpy(buf, f->contents + f->off, count)
    f->off += count;
    ...
}
```

- Users can extend/replace environment w/o any knowledge of KLEE's internals
  - Often the first part of KLEE users experiment with
- Users can choose precision
  - fail system calls? etc.
- Currently: effective support for symbolic command line arguments, files, links, pipes, ttys, environment vars

# Statistics

Core Engine

Stats

The screenshot displays the klee-last/run.istats application interface. The window title is "klee-last/run.istats [./echo.bc]". The interface includes a menu bar (File, View, Go, Settings, Help) and a toolbar with icons for Open, Reload, Force Dump, Up, Back, Forward, Show Relative Costs, Percentage Relative to Parent, Do Cycle Detection, and Instructions. The main area is divided into two panes:

- Flat Profile:** A table showing function call statistics. The columns are "Incl.", "Self", "Called", "Function", and "Location". The "user\_main" function is highlighted in blue.
- Source Code:** A view of the C source code for "echo.c". The line number 138 is highlighted, corresponding to the "setlocale" call in the Flat Profile.

Incl.	Self	Called	Function	Location
300 673	1	(0)	main	assembly.ll
300 672	544	1	_uClibc_main	assembly.ll: _uC
289 676	1 317	1	user_main	assembly.ll: echo
220 591	43	1	setlocale	assembly.ll: locale
217 627	692	1	newlocale	assembly.ll: locale
107 771	8 522	2	_locale_set_l	assembly.ll: locale
98 085	98 085	19	memcpy	assembly.ll: mem
95 083	2 083	2	init_cur_collate	assembly.ll: locale
75 006	16 306	20	getenv	assembly.ll: gete
56 650	56 650	866	memcmp	assembly.ll: mem
48 510	18 756	6	find_locale	assembly.ll: locale
35 009	500	25	exit	assembly.ll: atex
33 484	1 450	25	_exit_handler	assembly.ll: atex
32 034	375	25	close_stdout	assembly.ll: close
31 659	2 500	50	close_stream	assembly.ll: close
26 788	5 095	50	fclose	assembly.ll: fclose
26 412	26 412	335	strncmp	assembly.ll: strnc
24 566	24 566	55	memset	assembly.ll: mem
21 688	149	1	parse_long_options	assembly.ll: long-c
19 850	1 350	50	close	assembly.ll: fd.c
10 524	91	1	version_etc_va	assembly.ll: versio
10 026	59	1	rpl_vfprintf	assembly.ll: vfprin
9 720	664	1	vasnprintf	assembly.ll: vasnpr
8 362	15	1	_uClibc_init	assembly.ll: _uCli
7 555	37	1	snprintf	assembly.ll: snprin
7 518	101	1	vsnprintf	assembly.ll: vsnpr
7 417	123	1	fprintf	assembly.ll: vfpr
6 739	4	1	_locale_init	assembly.ll: locale
6 735	129	1	_locale_init_l	assembly.ll: locale
6 549	79	1	getopt_long	assembly.ll: getopt
6 470	199	1	getopt_internal	assembly.ll: getopt
6 271	1 095	1	getopt_internal_r	assembly.ll: getopt
5 618	1 189	20	fwrite_unlocked	assembly.ll: fwrite
4 970	2 660	35	_stdio_WRITE	assembly.ll: _WRI
4 894	236	1	_ppfs_init	assembly.ll: _vpri
4 611	4 611	62	strlen	assembly.ll: strlen

```
# | Source ('/home/ddunbar/public/cu_test/coreutils-6.11/obj-llvm/src/././src/echo.c')
116 | }
117 | }
118 |
119 | /* Print the words in LIST to standard output. If the first word is
120 | '-n', then don't print a trailing newline. We also support the
121 | echo syntax from Version 9 unix systems. */
122 |
123 | int
124 | main (int argc, char **argv)
125 | {
126 |     34 bool display_return = true;
127 |     723 1 call to 'klee_init_env' (assembly.ll: klee_init_env.c)
128 |     bool allow_options =
129 |         (! getenv ("POSIXLY_CORRECT"));
130 |     11 || (! DEFAULT_ECHO_TO_XPG && 1 < argc && STREQ (argv[1], "-n"));
131 |     4 215 1 call to 'getenv' (assembly.ll: getenv.c)
132 |
133 |     /* System V machines already have a /bin/sh with a v9 behavior.
134 |        Use the identical behavior for these machines so that the
135 |        existing system shell scripts won't barf. */
136 |     2 bool do_v9 = DEFAULT_ECHO_TO_XPG;
137 |
138 |     initialize_main (&argc, &argv);
139 |     5 program_name = argv[0];
140 |     138 2 setlocale (LC_ALL, "");
141 |     220 591 1 call to 'setlocale' (assembly.ll: locale.c)
142 |     bindtextdomain (PACKAGE, LOCALEDIR);
143 |     textdomain (PACKAGE);
144 |
145 |     2 atexit (close_stdout);
146 |     136 1 call to 'atexit' (assembly.ll: _atexit.c)
147 |
148 |     4 if (allow_options)
149 |         11 parse_long_options (argc, argv, PROGRAM_NAME, PACKAGE_NAME, VERSION,
150 |                                21 688 1 call to 'parse_long_options' (assembly.ll: long-options.c)
151 |                                usage, AUTHORS, (char const *) NULL);
```

Good support for producing and visualizing a variety of statistics, associated with different entities and events

# Non-determinism in SymEx and KLEE


---

- Any good experiment needs to take non-determinism into account
- Sources of non-determinism include constraint solving, search heuristics, LLVM versions, memory allocation
  - Currently fixing implementation-level non-determinism, such as hash tables indexed by memory addresses, which can differ across runs


# Example: Constraint solving optimization in KLEE

Approach: run baseline KLEE for 30', rerun in the same configuration with optimizations

30 minutes



10 minutes



# Example 2: Coverage optimization in KLEE

---

Approach: take same benchmarks from paper X,  
rerun KLEE with coverage optimization

Baseline (LLVM 2.3)

60% coverage

Baseline (LLVM 3.4)

80% coverage

Optimized (LLVM 3.4)

80% coverage

# Dynamic Symbolic Execution

---

- Program analysis technique that can be use to automatically explore paths through a program
- Can generate inputs achieving high-coverage and exposing bugs in complex software

# KLEE: Freely Available as Open-Source

<http://klee.github.io/>

- Popular symbolic execution tool with an active user and developer base
- Extended in many interesting ways by several groups from academia and industry, in areas such as:
  - exploit generation
  - wireless sensor networks/distributed systems
  - automated debugging
  - client-behavior verification in online gaming
  - GPU testing and verification
  - etc. etc.