

# Modern Symbolic Execution

Austin Cory Bart

CS-6304 Program Analysis

11/10/2015

# Papers

- Cadar et al, “Symbolic Execution for Software Testing in Practice – a Preliminary Assessment”, ICSE 2011.
- C. Cadar & K. Sen, “Symbolic Execution for Software Testing: Three Decades Later”, CACM, Feb 2013, p 82-90



Cristian Cadar  
PhD from Stanford  
Now at Imperial College London  
KLEE, EXE



Koushik Sen:  
PhD from University of Illinois at Urbana-Champaign  
Now at UC Berkley  
DART, Latest, CUTE, jCUTE, Jalangi

# Outline

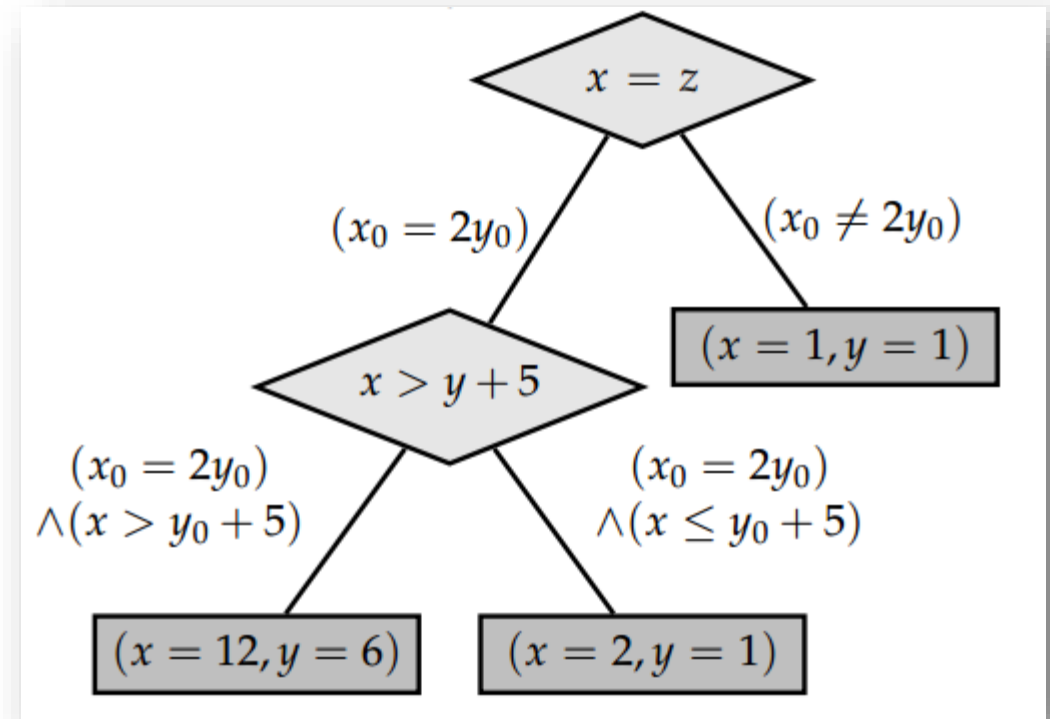
- Motivation
- Symbolic Execution Techniques
  - EGT
  - Concolic Testing
- Challenges
  - Path Explosion
  - Constraint Solving
  - Concurrency
- Tools

# Why Care?

- Automatic Software testing
- Systems and above

# Execution Tree

```
1  foo(int x, int y){
2      z = 2*y;
3      if (x == z){
4          if (x > y + 5){
5              //some error
6          }
7      }
8  }
```



# Concolic vs. EGT

## **Concolic**

- Simultaneous Concrete and Symbolic
- Needs initial concrete values
- Multiple runs

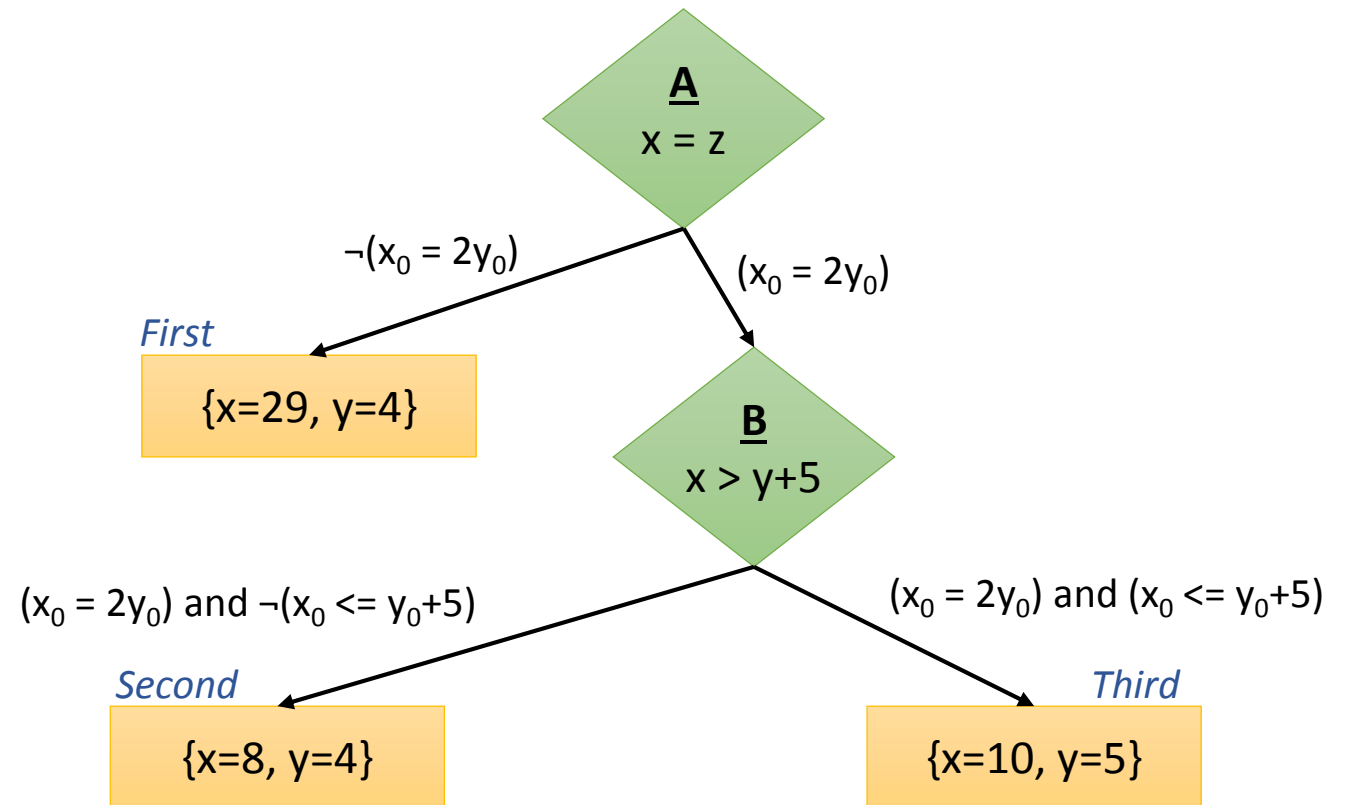
## **EGT**

- Concrete values generated “on-demand”
- No initial concrete values
- Forking execution
- More similar to vanilla Symbolic Execution

# Concolic Example #1

```
1  foo(int x,int y){
2      z = 2*y;
3      if (x == z){
4          if (x > y + 5){
5              //some error
6          }
7      }
8  }
```

- Initial random input:  $\{x=29, y=4\}$
- First run, A[false]:
  - $x_0 \neq 2y_0$
  - Negate conjunct, so  $x_0 == 2y_0$
  - New input:  $\{x=8, y=4\}$
- Second Run, A[true] and B[false]:
  - $(x_0 == 2y_0)$  and  $(x_0 \leq y_0+5)$
  - Negate *new* conjunct, so  $(x_0 > y_0+5)$
  - New input:  $\{x=10, y=5\}$
- Third Run, A[true] and B[True]:
  - $(x_0 == 2y_0)$  and  $(x_0 > y_0+5)$
  - All conjuncts tested. Complete.



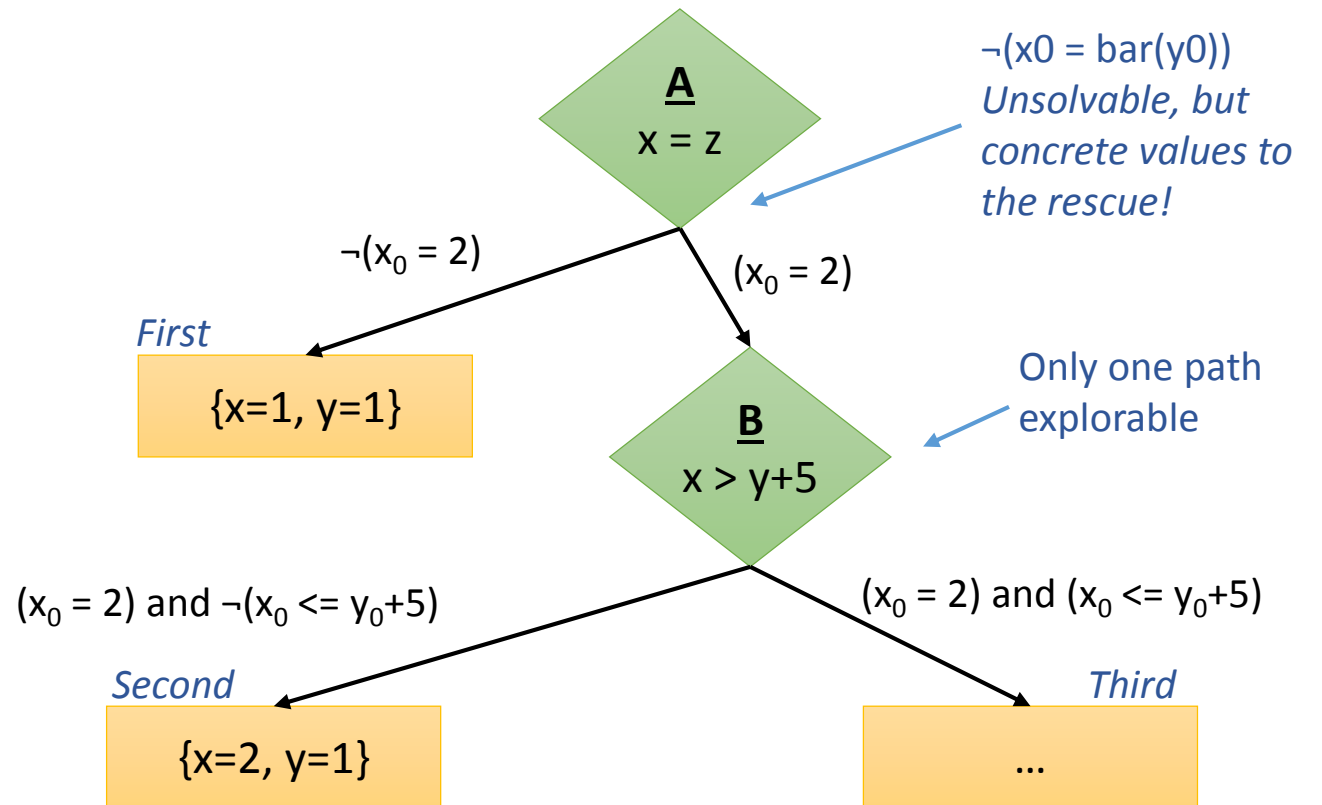
# Concolic Example #2

- Initial random input:  $\{x=1, y=1\}$
- First run, **Concrete Evaluation!**, A[false]:
  - $x_0 \neq 2$
  - Negate conjunct, so  $x_0 == 2$
  - New input:  $\{x=2, y=1\}$
- ...

```
1 foo(int x, int y){
2   z = bar(y)
3   if (x == z){
4     if (x > y + 5){
5       //some error
6     }
7   }
8 }
```

```
bar(int w){
  return 2*w;
}
```

(Executable, but source code not available)

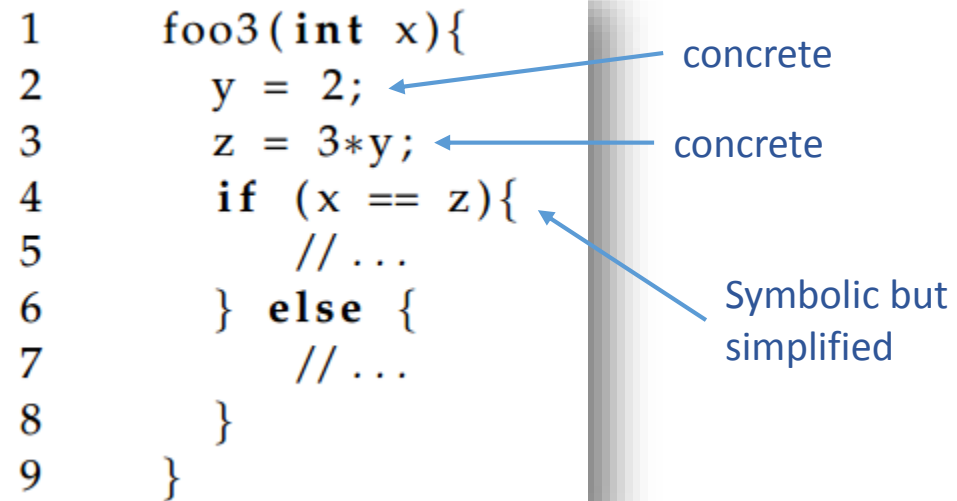




# EGT Example

```
1  foo3(int x){  
2    y = 2; ← concrete  
3    z = 3*y; ← concrete  
4    if (x == z){  
5        // ...  
6    } else {  
7        // ...  
8    }  
9 }
```

Symbolic but simplified

The diagram shows a code block with three blue arrows pointing from labels on the right to specific lines of code. The label 'concrete' appears twice, with arrows pointing to lines 2 and 3. The label 'Symbolic but simplified' has an arrow pointing to line 4. The code block is enclosed in a white box with a grey shadow.

# Commonalities

- Overcome
  - External code
  - Hardware imprecision (e.g., floating points)
  - Constraint Solver timeouts
- Sound, but not complete
- Automatic

# Key Challenges

And their solutions

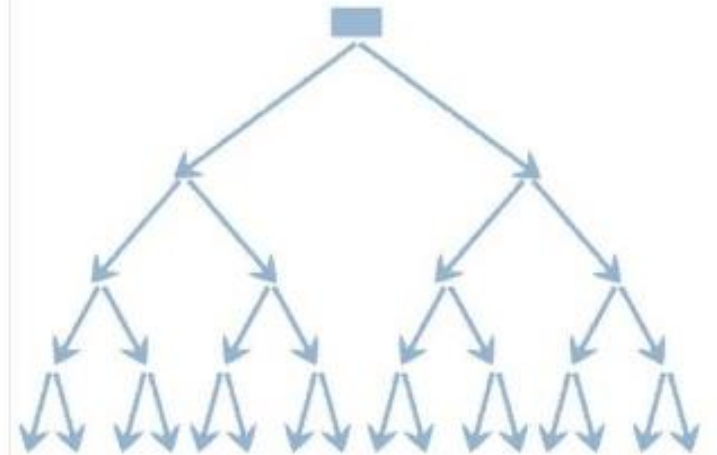
# Problem 1: Path Explosion

```
void process(char input[3]) {  
    int counter = 0;  
    if (input[0] == 'a') counter++;  
    if (input[1] == 'b') counter++;  
    if (input[2] == 'c') counter++;  
    if (counter >= 3) success();  
    error();  
}
```

- Exponentially many execution paths



4 conditional nodes



16 ( $2^4$ ) execution paths

# Infinite Execution Paths

```
1  void testme_inf () {  
2      int sum = 0;  
3      int N = sym_input();  
4      while (N > 0) {  
5          sum = sum + N;  
6          N = sym_input();  
7      }  
8  }
```

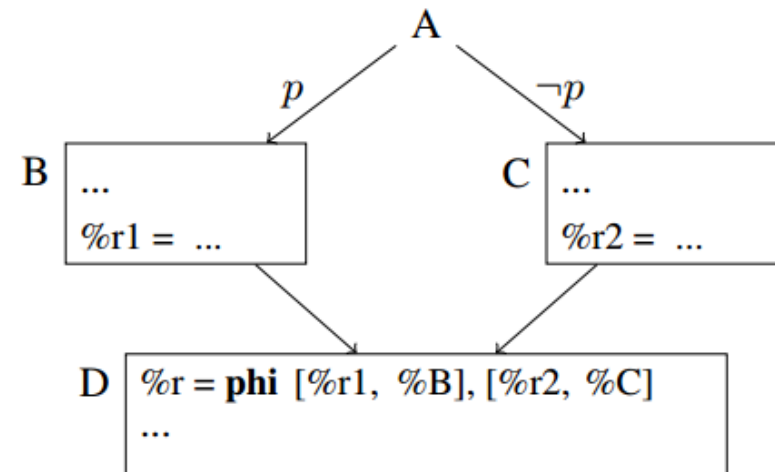
---

**Figure 3.** Simple example to illustrate infinite number of execution paths.



# Solution 1.2: Select Statements

- Merge If-conditions into Select statements
- Phi-node folding (if-conversion)
  - Static-single assignment (SSA)
  - Diamond-shaped if statements
  - Unconditionally execute and select result
  - Side-effects can occur!
- Passes the buck to the Constraint Solver



**Figure 3.** Diamond control flow pattern.

# Solutions 1.3, 1.4, 1.5: Other techniques

- Cache and reuse the analysis of lower-level functions
  - Pre-/post- condition summaries
- Lazy Test Generation
  - “The technique first explores, using dynamic symbolic execution, an abstraction of the function under test by replacing each called function with an unconstrained input.”
    - Strlen becomes a symbolic input that can represent any integer
- Prune redundant paths
  - Redundancy: same program path, same symbolic constraints



# Problem 2: Constraint Solving

- NP Complete (although practical in practice)
- Dominates the runtime

# Solution 2.1: Irrelevant Constraint Elimination

$$(x + y > 10) \wedge (z > 0) \wedge (y < 12) \wedge (z - x = 0)$$



*Negated conjunct  
for new inputs*

$$(x + y > 10) \wedge (z > 0) \wedge \neg(y < 12)$$



*No relationship*

$$(x + y > 10) \wedge \neg(y < 12)$$

# Solution 2.2: Incremental Solving

- Cached constraint solutions

$$(x + y < 10) \wedge (x > 5) \Rightarrow \{x = 6, y = 3\}$$


- Two situations:
  - Subset of a cached constraint: Easy, use the cached inputs!
  - Superset of a cached constraint: Test the inputs!

$$(x + y < 10) \wedge (x > 5) \wedge (y \geq 0)$$

- *“In practice, adding constraints often does not invalidate an existing solution”*

# Problem 3: Memory Modelling

- 32-bit integer



- Pointers
  - `b[7]` vs. `b[i]` vs. `a[b[i]]`

# Problem 4: Handling Concurrency

- Complex Data Inputs
- Distributed Systems
- GPGPU Programs
  
- Race conditions cause interleaving explosion

# Solution 4.1: “Race Detection and Flipping Algorithm”

- Adaption of Concolic by Koushik Sen and Gul Agha
- Identify identical interleavings
  - Race conditions are collected during execution alongside path constraints
  - Race for two events if:
    - Stem from different threads
    - Both access the same memory location without locks
    - Order permutable by changing thread scheduling
  - Sequence of Triples: (thread, label, shared memory access type)
  - Types of race conditions: sequential, shared-memory access precedence, causal and race relation.
- Works by varying execution times
- Vector clocks (integer vectors) to record thread execution

# Tool Rundown

- DART: First concolic testing (C)
- CUTE: Multi-threaded DART, dynamic data structures (C)
- jCUTE: CUTE for Java
- CREST: Concolic testing for experimenting with heuristics (C)
- EXE: EGT approach for Bit-level accuracy using STP (C)
- KLEE: Concurrent states, external data, heavily extended (LLVM)
- SAGE: Microsoft Windows, uses fuzzing (x86 binaries)
- PEX: Microsoft, focuses on more pure symbolic (.NET)

# Conclusions

- Mixing symbolic and concrete is useful
- Many successes
- But still a lot left to do
  - Parallel
  - Constraint Solving
  - Memory models
  - Heuristics