

Chopping Code for More Modular and Scalable Symbolic Execution

Cristian Cadar

Department of Computing
Imperial College London



Joint work with

David Trabish (Tel Aviv University)
Andrea Mattavelli (Imperial College London)
Noam Rinetzky (Tel Aviv University)

Symbolic Execution or Dynamic Symbolic Execution (DSE)

Program analysis technique for ***automatically exploring paths*** through a program

Applications in:

- Bug finding
- Test generation
- Vulnerability detection and exploitation
- Equivalence checking
- Debugging
- Program repair
- etc. etc.



Modern Symbolic Execution

review articles



- Symbolic execution introduced in 1970s
 - Boyer, Elspass, Levitt (SRI)
 - Clarke (UMass Amherst)
 - King (IBM Research)
- Revived in 2005 in the form of *dynamic symbolic execution*:
 - DART system (Bell Labs)
 - EGT system (Stanford)

Dynamic Symbolic Execution

PyExZ3

SAGE

Jalangi2

angr



Pex

CATG

SymDroid



CiVL



KLOVER

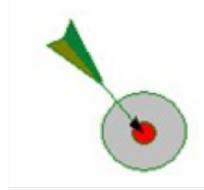
Otter

BinSE

Symbolic
PathFinder



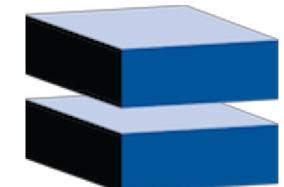
PathGrind



DART

LDSE

JDart



S²E

jCUTE

SymJS

Miasm

CUTE

Kite





Webpage: <http://klee.github.io/>
Code: <https://github.com/klee>
Web version: <http://klee.doc.ic.ac.uk/>

- Active user and developer base with over 300 subscribers on the mailing list and over 50 contributors listed on GitHub
- Academic impact:
 - Over 2K citations to original KLEE paper
 - From many different research communities: testing, verification, systems, software engineering, security, etc.
 - Several influential systems using KLEE: AEG, Angelix , BugRedux , Cloud9, GKLEE, KleeNet, KLEE-UC, S2E, SemFix, etc.
- Growing impact in industry:
 - **Fujitsu**: [PPoPP 2012], [CAV 2013], [ICST 2015], [IEEE Software 2017], [KLEE Workshop 2018], **Hitachi**: [CPSNA 2014], [ISPA 2015], [EUC 2016], **Trail of Bits**: <https://blog.trailofbits.com/>, **Intel**: [WOOT 2015], **NASA Ames**: [NFM 2014], **Baidu**: [KLEE Workshop 2018]

From Whole-Program Analysis... ...To More Localized Tasks

- Most work on modern symbolic execution:
 - Whole-program test generation
 - Whole-program bug-finding
- More recently attention shifted to more localized tasks:
 - Patch testing
 - Debugging
 - Bug reproduction
 - Program repair
 - etc.
- Which one is easier?

Opportunity of more localized tasks:
Prune a large part of the search space

Chopped Symbolic Execution

- Some code fragments are unrelated to certain tasks
 - But symbolic execution can spend lots of time unnecessarily analyzing them
- Determining precisely if a part of the code is unrelated is hard
 - Often, most computation in a code fragment is unrelated, but not all

IDEA:

- 1) Guess unrelated code fragments (manually or via lightweight analysis)
- 2) Speculatively skip these code fragments
- 3) If their side effects are ever needed, execute relevant paths only

Chopped Symbolic Execution

```
int j; // symbolic
int k; // symbolic
int x = 0;
int y = 0;
```

```
void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    }
    else target2;
}
```

```
void f() {
    if (k > 0)
        x = 1;
    else
        if (j > 0)
            y = 1;
        else
            y = 0;
}
```



Chopped Symbolic Execution

```
int j; // symbolic  
int k; // symbolic  
int x = 0;  
int y = 0;
```

```
void main() {  
    f();  
    if (j > 0) {  
        if (y)  
            target1;  
    }  
    else target2;  
}
```

$\text{Ref}(\text{main}) = \{j, y\}$

```
void f() {  
    if (k > 0)  
        x = 1;  
    else  
        if (j > 0)  
            y = 1;  
        else  
            y = 0;  
}
```

$\text{Mod}(f) = \{x, y\}$

Chopped Symbolic Execution

```
int j; // symbolic
int k; // symbolic
int x = 0;
int y = 0;
```

```
void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    }
    else target2;
}
```

```
void f() {
    if (k > 0)
        x = 1;
    else
        if (j > 0)
            y = 1;
        else
            y = 0;
}
```

Dependent load

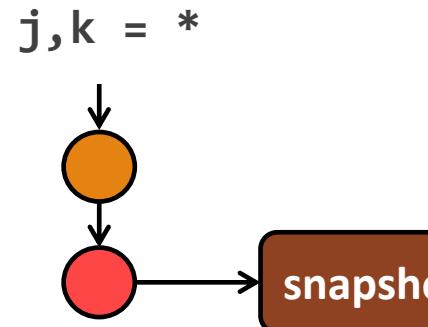
Chopped Symbolic Execution

```
j,k = *
```



```
void main() {  
    f();  
    if (j > 0) {  
        if (y)  
            target1;  
    }  
    else target2;  
}
```

Chopped Symbolic Execution

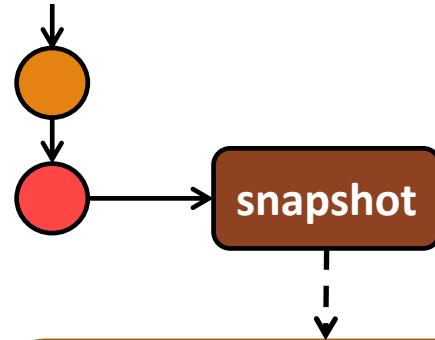


→

```
void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    }
    else target2;
}
```

Chopped Symbolic Execution

j,k = *

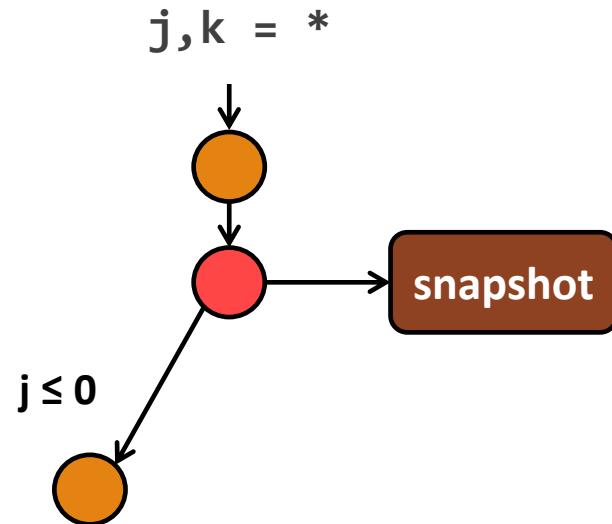


Program counter: *line 2*
Stack = [main]
Path constraints: {}
Memory: {x = 0, y = 0, k = ...}



```
void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    }
    else target2;
}
```

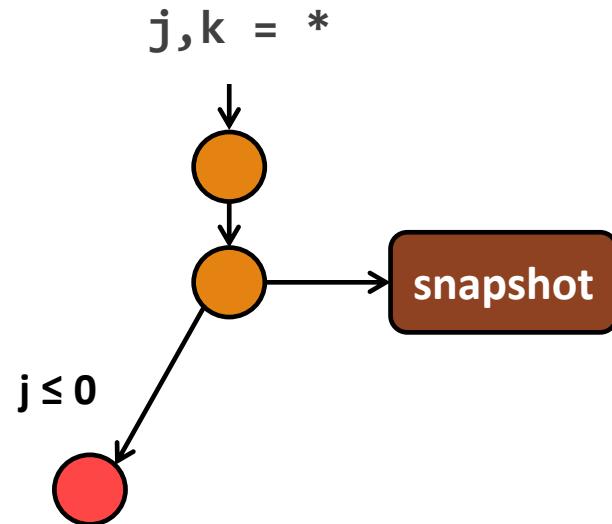
Chopped Symbolic Execution



→

```
void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    }
    else target2;
}
```

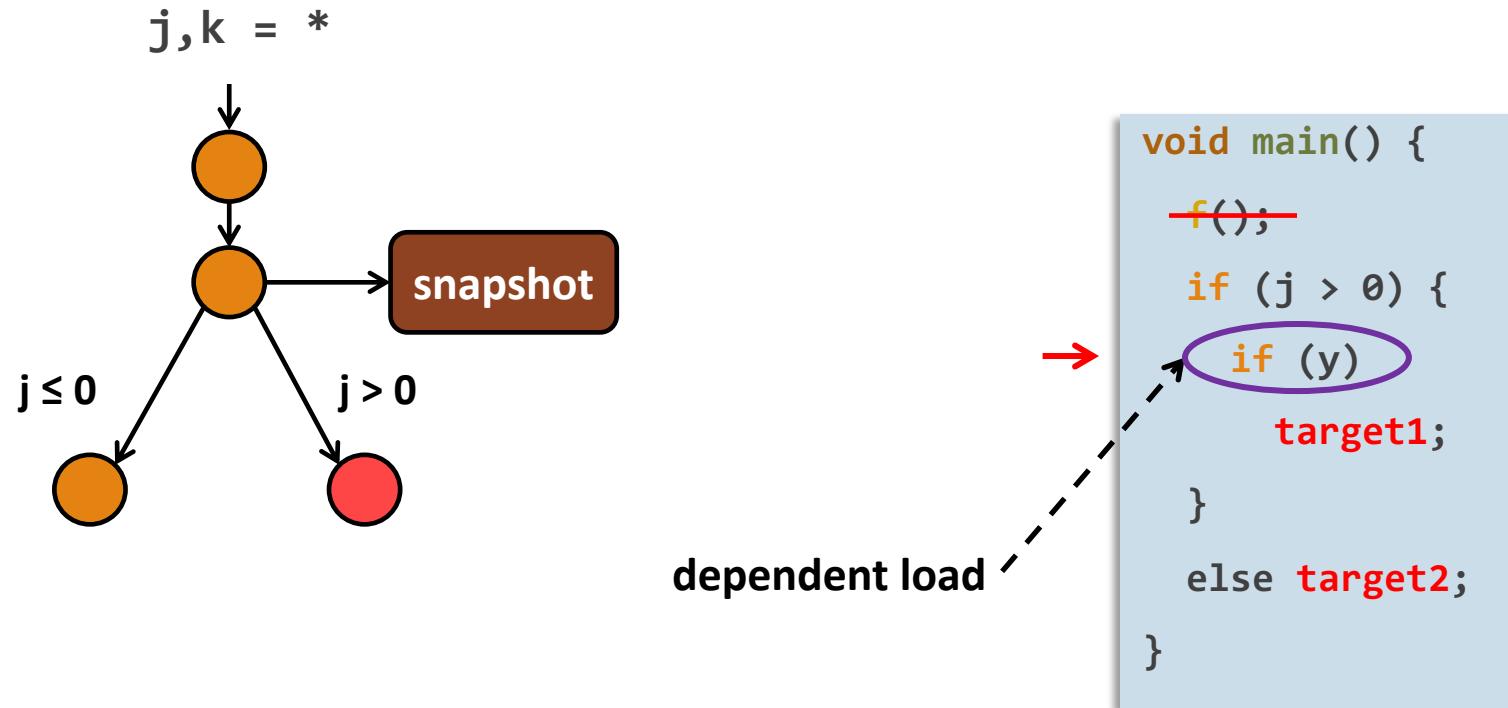
Chopped Symbolic Execution



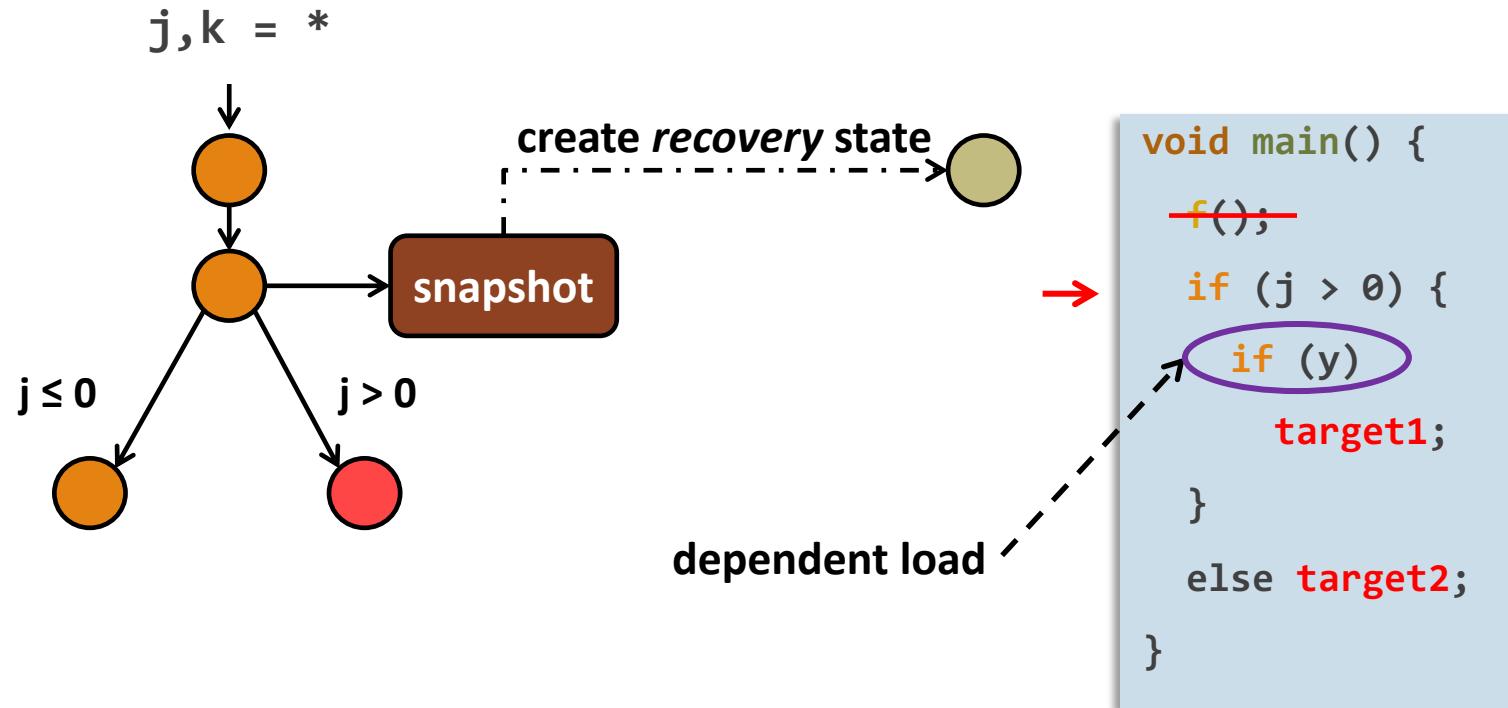
```
void main() {  
    f();  
    if (j > 0) {  
        if (y)  
            target1;  
    }  
    else target2;  
}
```

A red arrow points from the "snapshot" box to the code snippet, indicating that the snapshot taken during symbolic execution corresponds to the state of the program at that point.

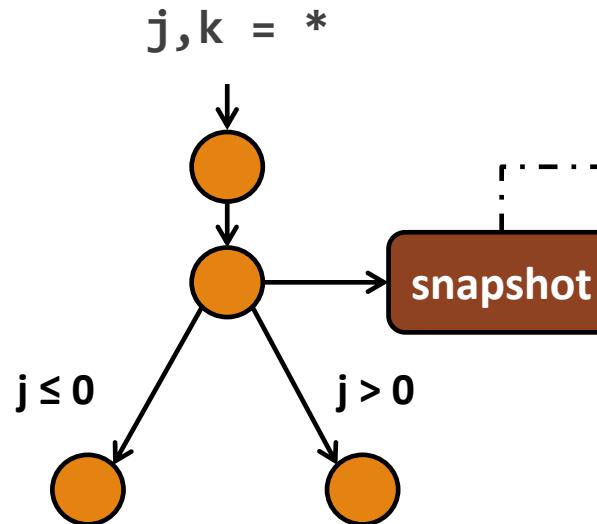
Chopped Symbolic Execution



Chopped Symbolic Execution



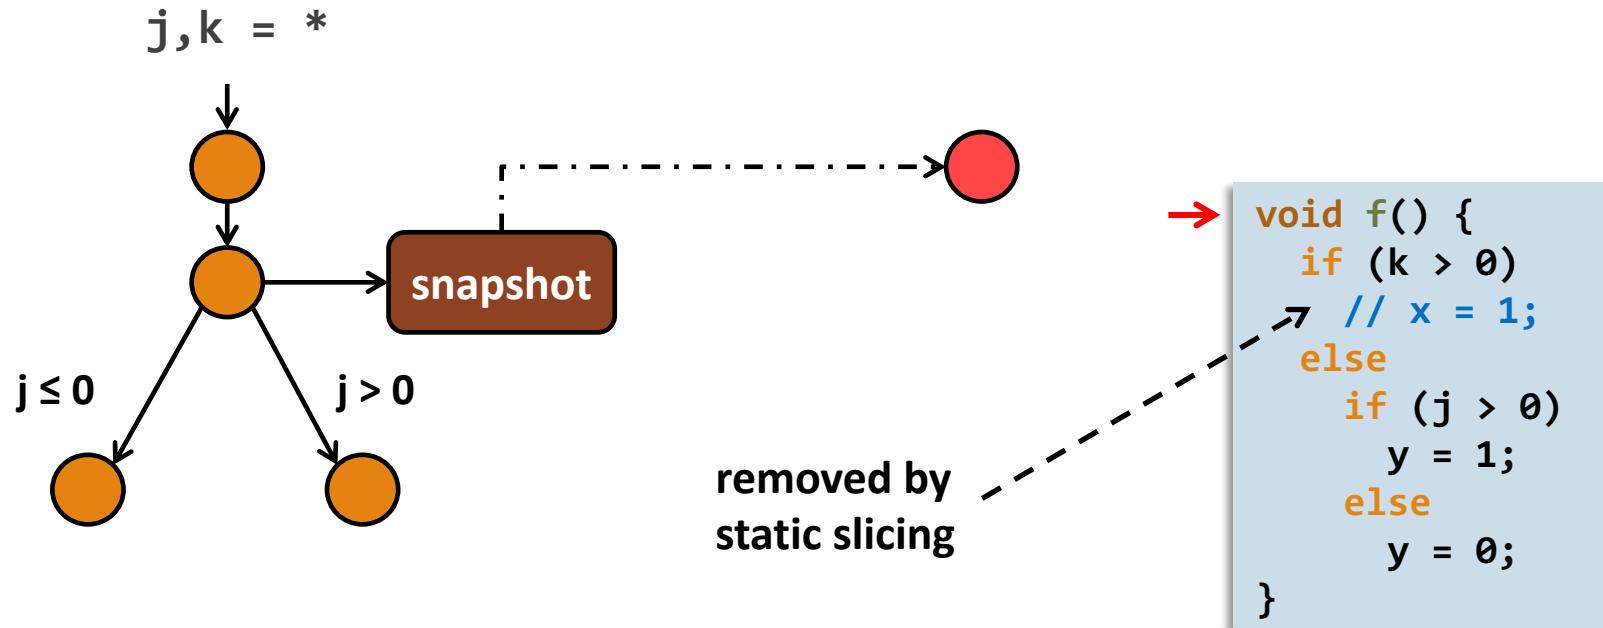
Chopped Symbolic Execution



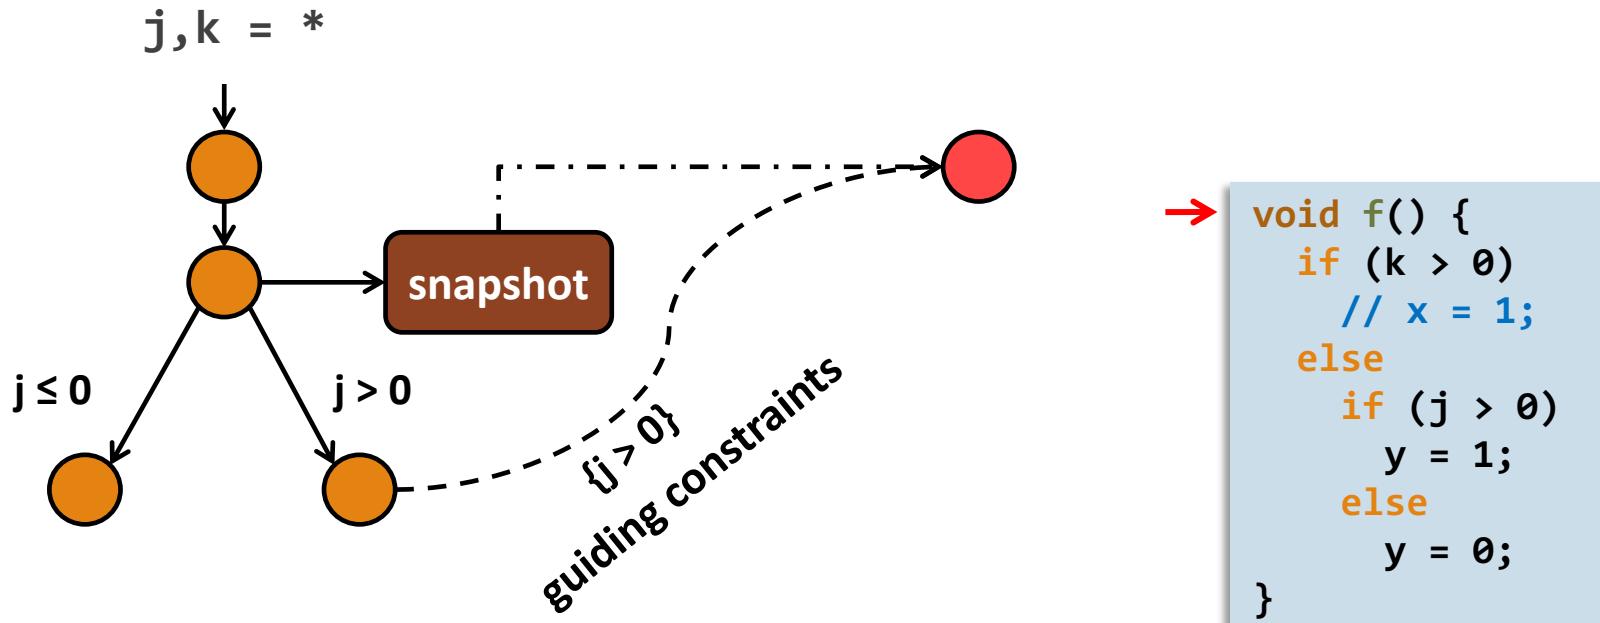
→

```
void f() {  
    if (k > 0)  
        x = 1;  
    else  
        if (j > 0)  
            y = 1;  
        else  
            y = 0;  
}
```

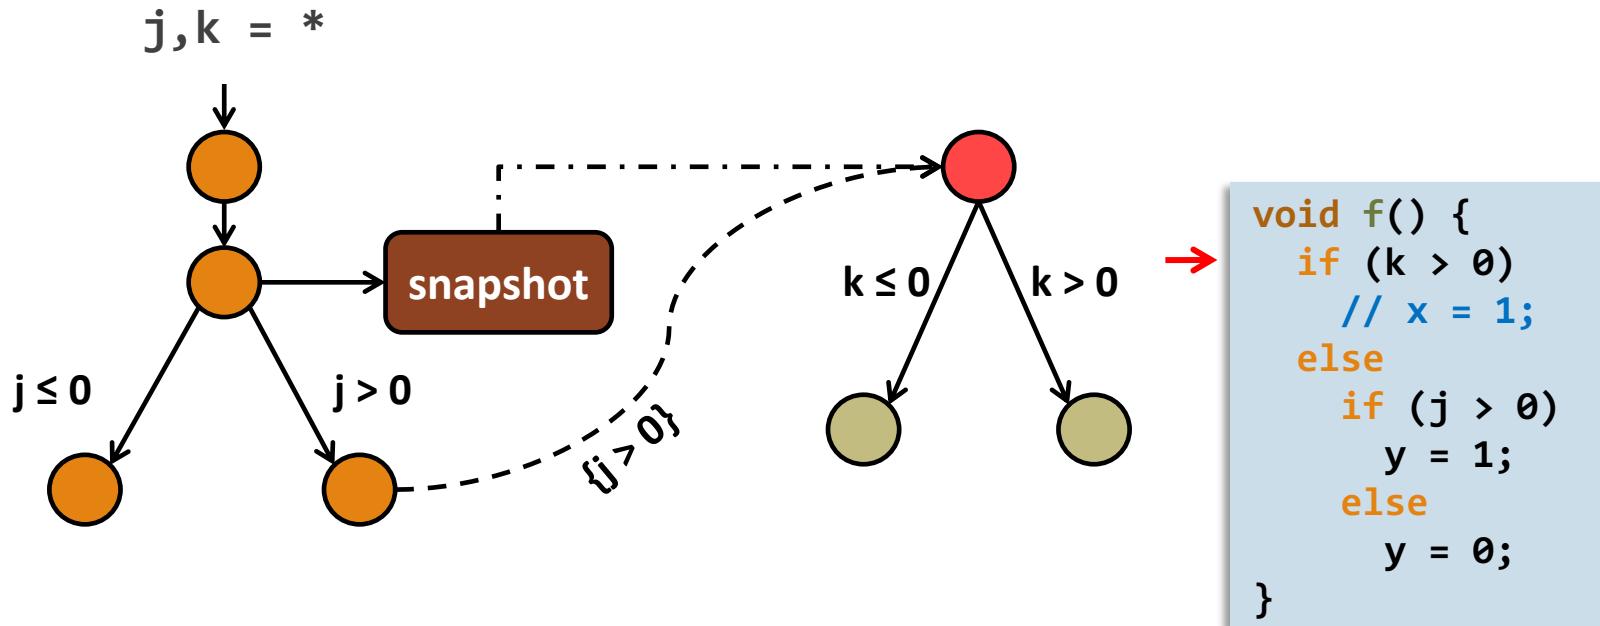
Chopped Symbolic Execution



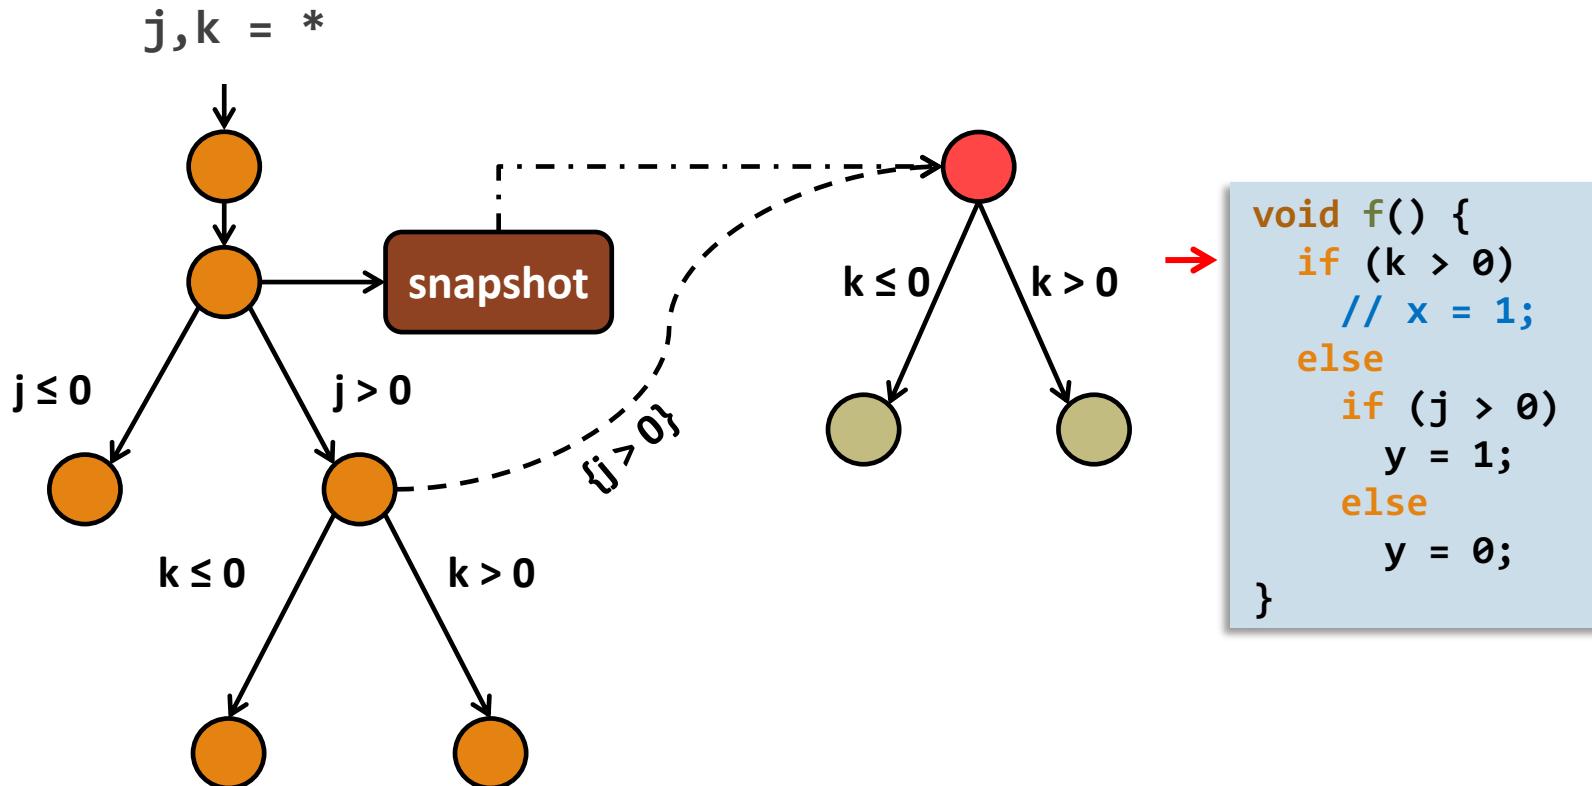
Chopped Symbolic Execution



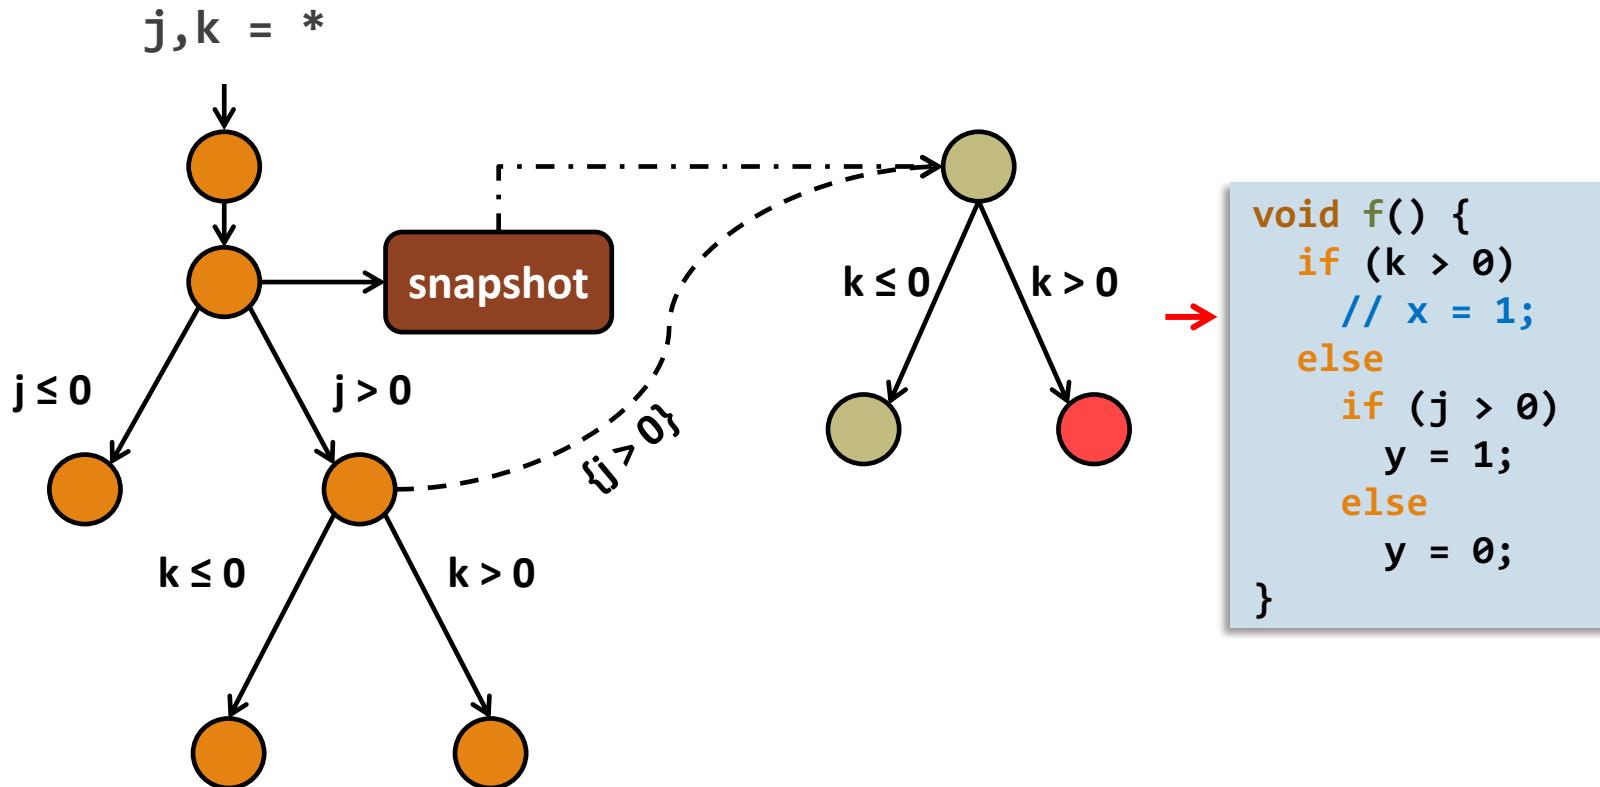
Chopped Symbolic Execution



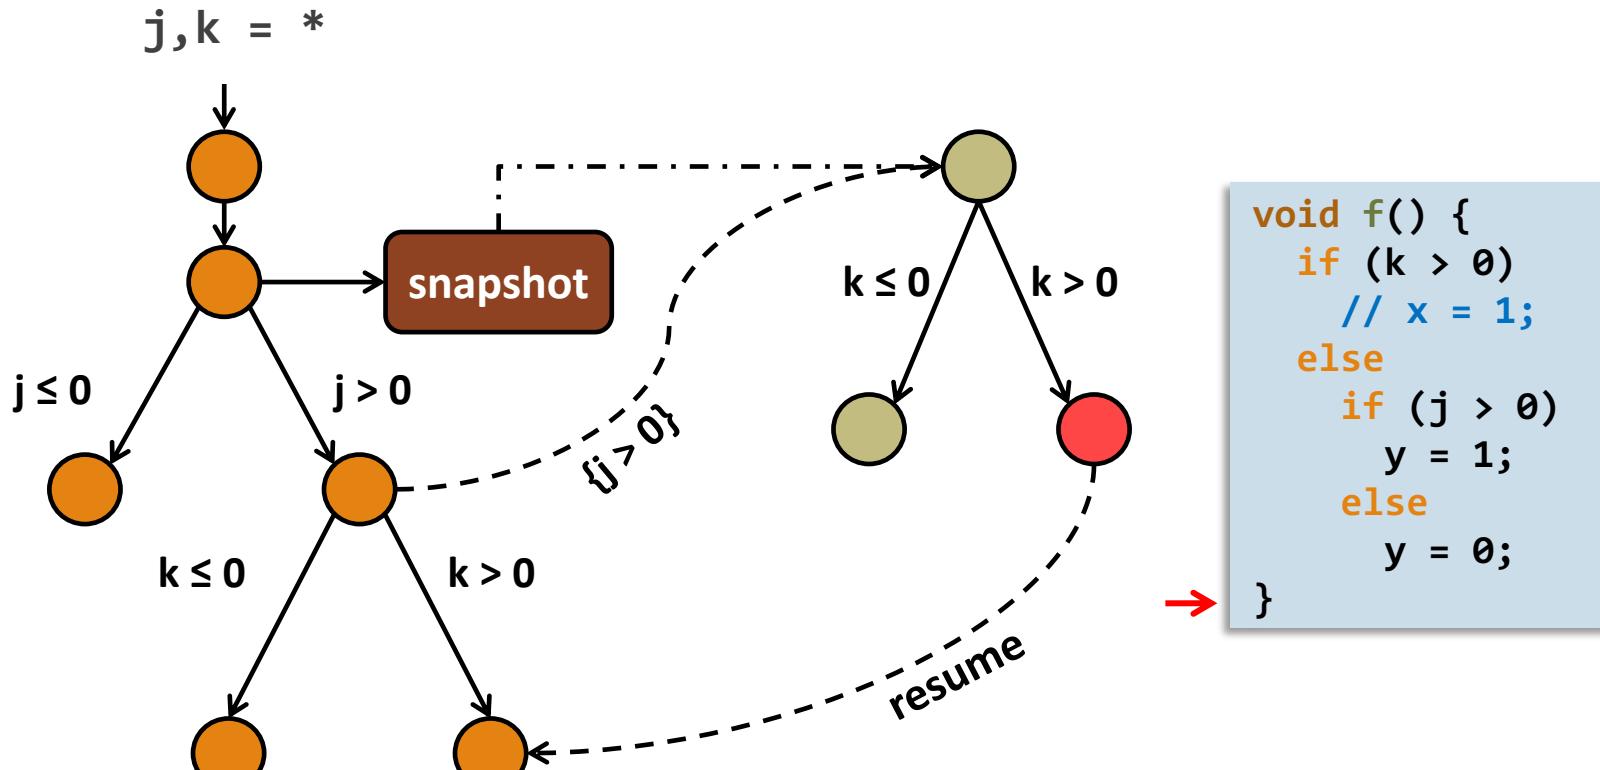
Chopped Symbolic Execution



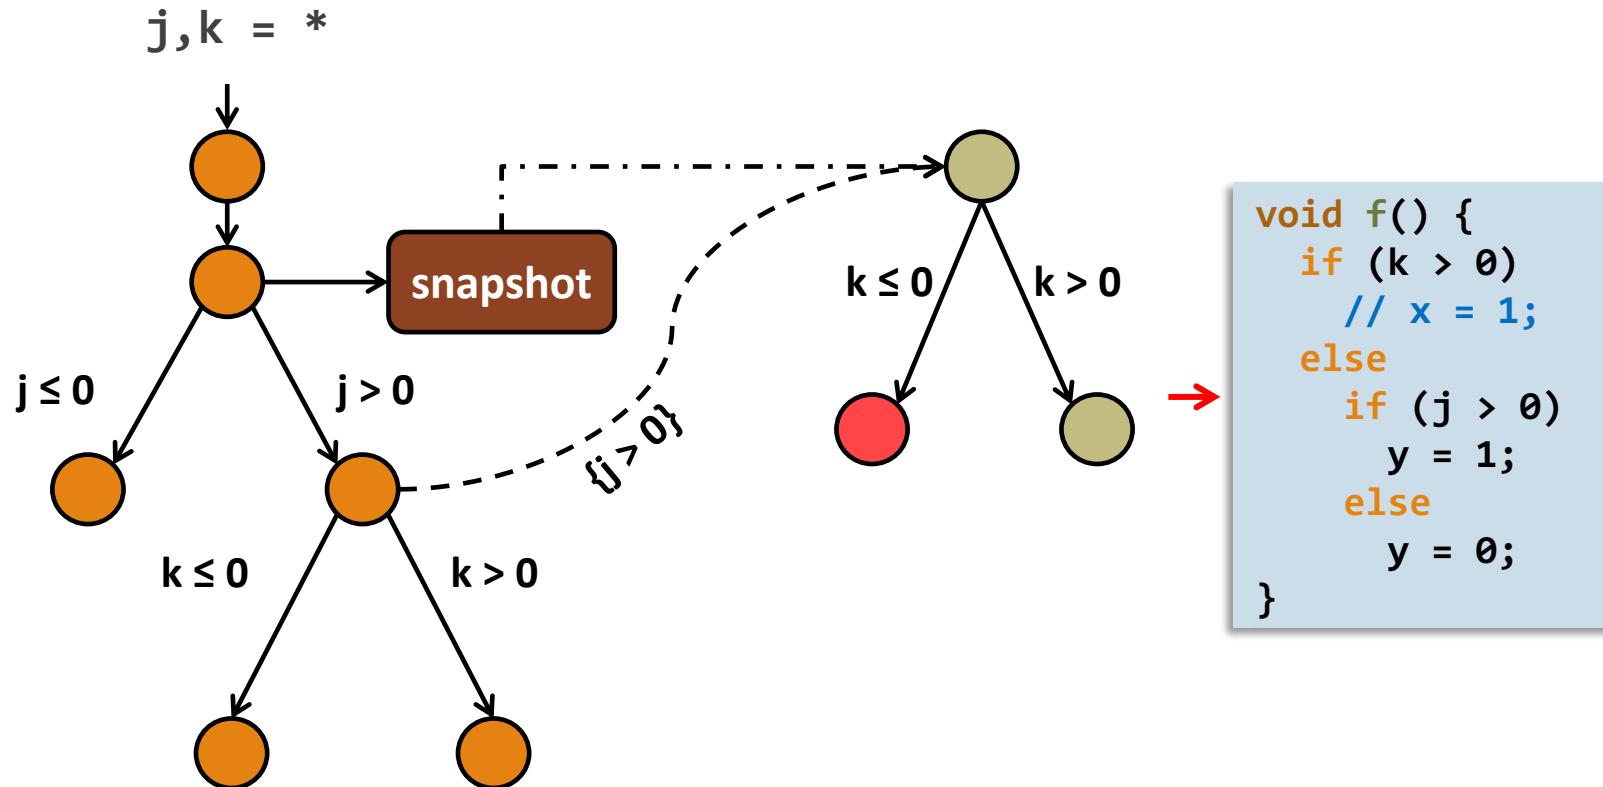
Chopped Symbolic Execution



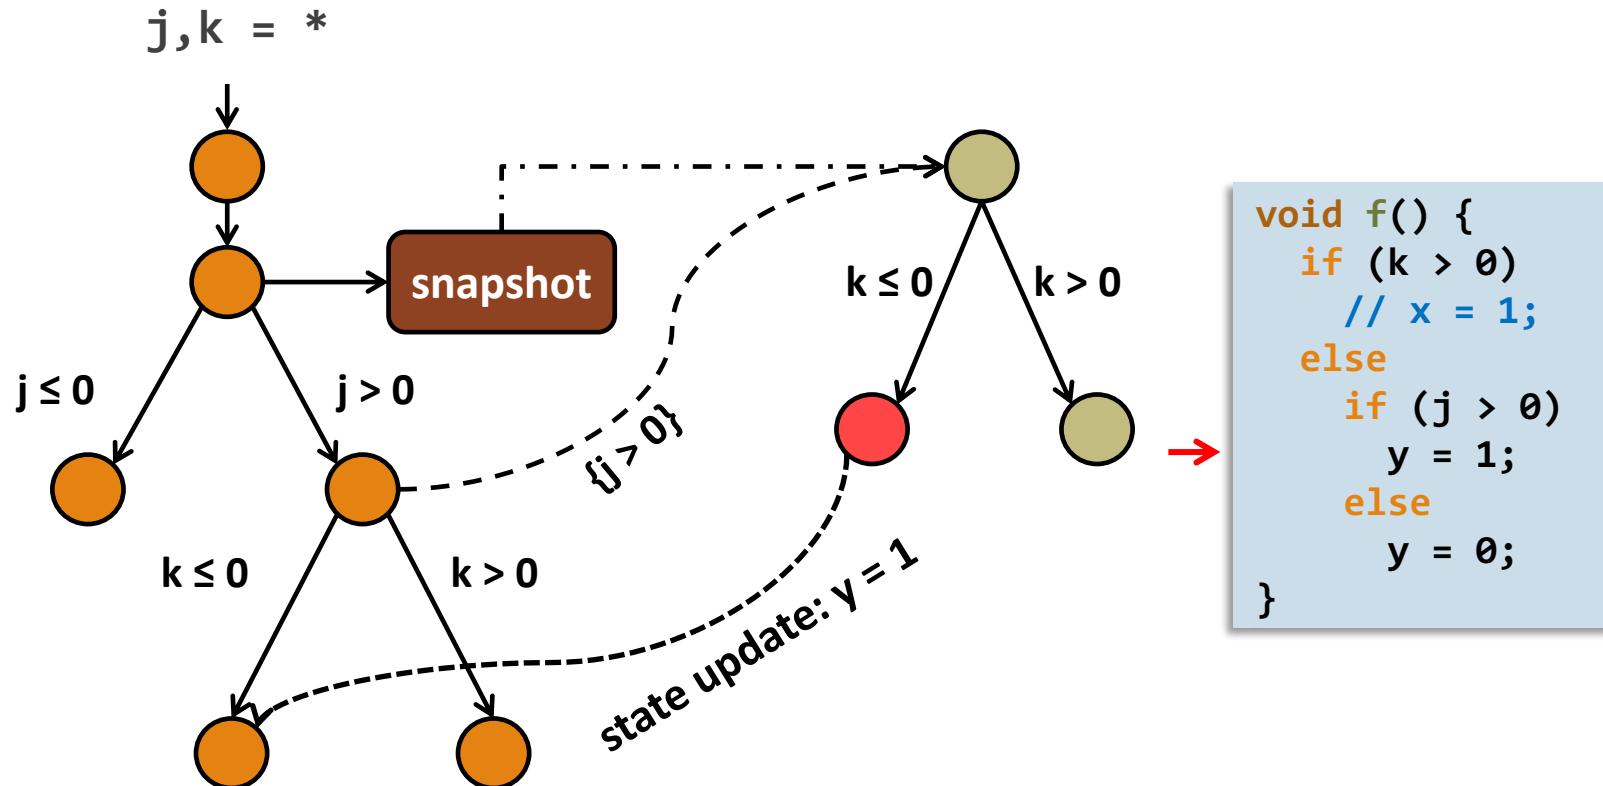
Chopped Symbolic Execution



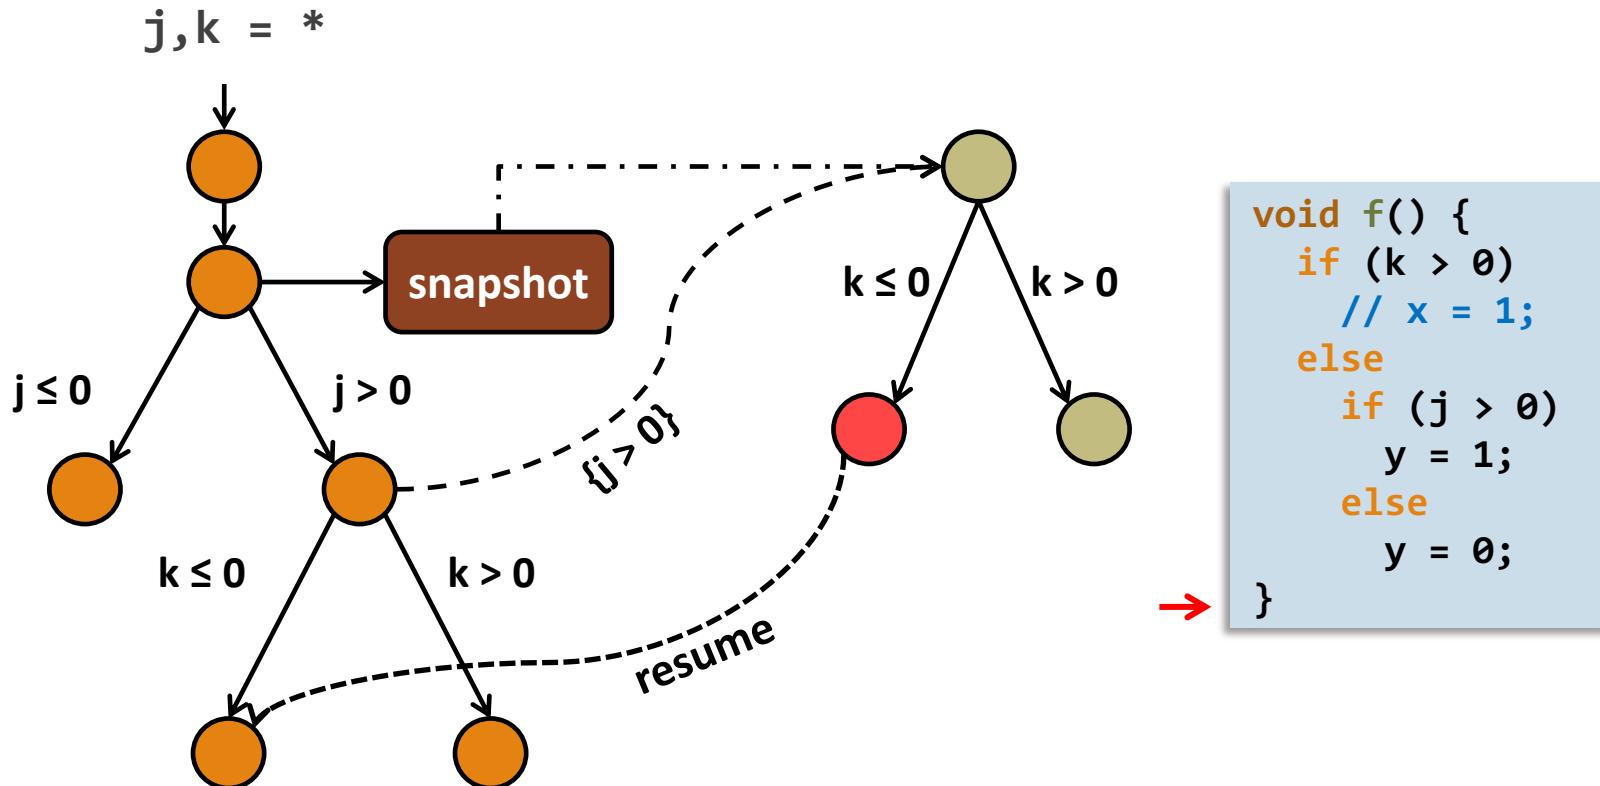
Chopped Symbolic Execution



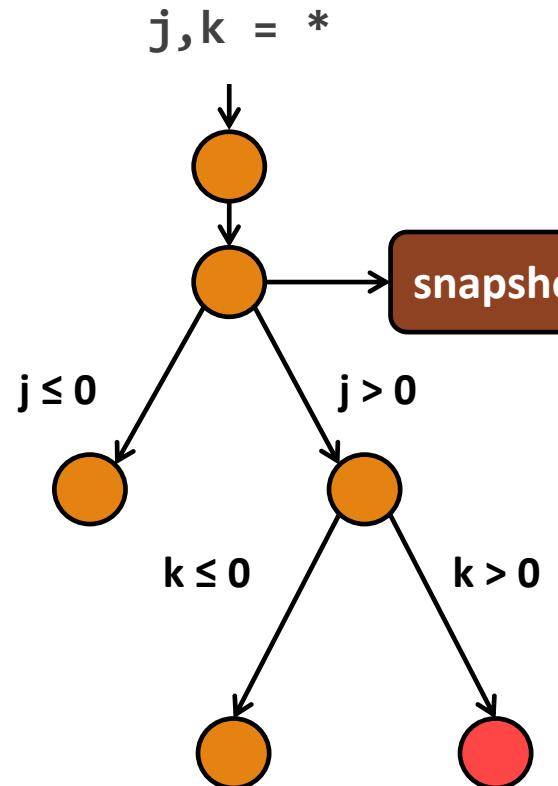
Chopped Symbolic Execution



Chopped Symbolic Execution



Chopped Symbolic Execution



```
void main() {  
    f();  
    if (j > 0) {  
        if (y)  
            target1;  
    }  
    else target2;  
}
```

```

void main() {
    f();
    if (j > 0) {
        if (y)
            bug1();
    }
    else bug2();
}

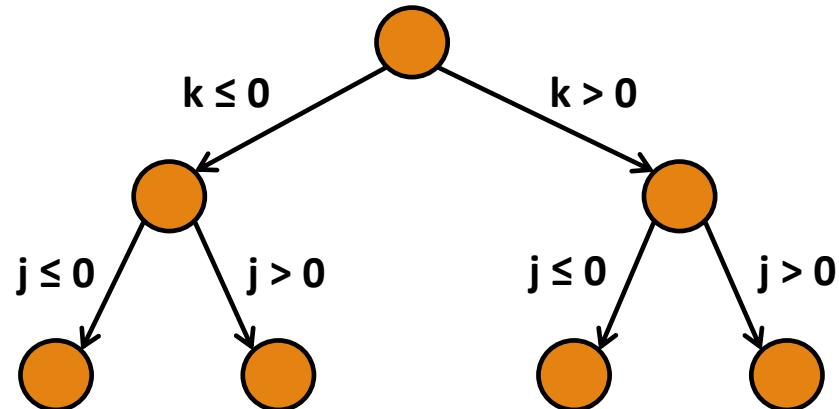
```

```

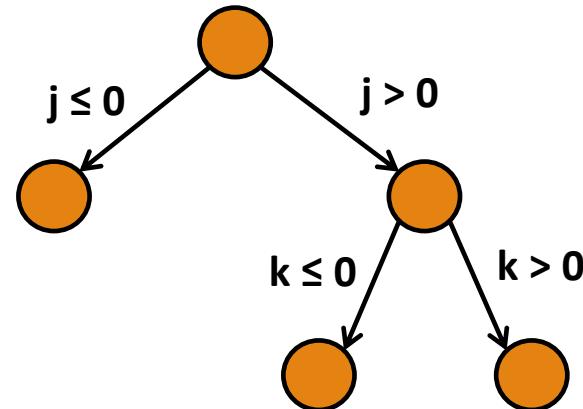
void f() {
    if (k > 0)
        x = 1;
    else
        if (j > 0)
            y = 1;
        else
            y = 0;
}

```

Execution Trees



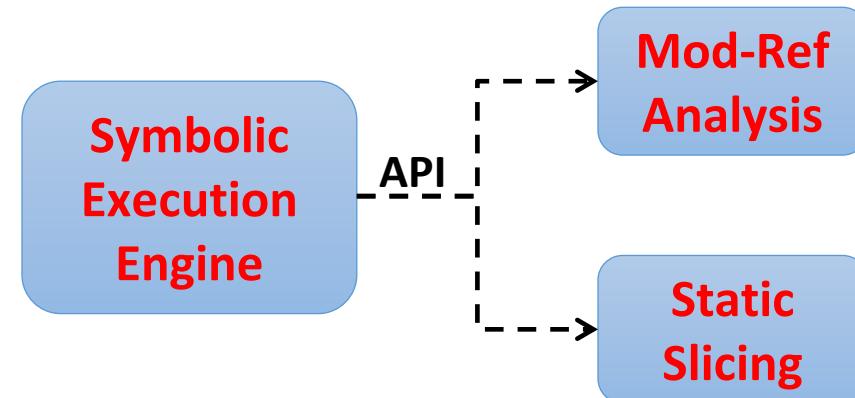
Standard SE



Chopped SE

Implementation: Chopper

- Symbolic execution based on KLEE
- Mod-ref analysis based on SVF
- Static slicing based on DG



Experiments

SECURITY VULNERABILITY REPRODUCTION

TEST SUITE AUGMENTATION

PATCH TESTING

```
address = optimizer.optimizeExpr(address, true);
StatePair zeroPointer = fork(state, Expr::createIsZero(address), true);
if (zeroPointer.first) {
    if (target)
        bindLocal(target, *zeroPointer.first, Expr::createPointer(0));
}
if (zeroPointer.second) { // address != 0
    ExactResolutionList rl;
    resolveExact(*zeroPointer.second, address, rl, "free");

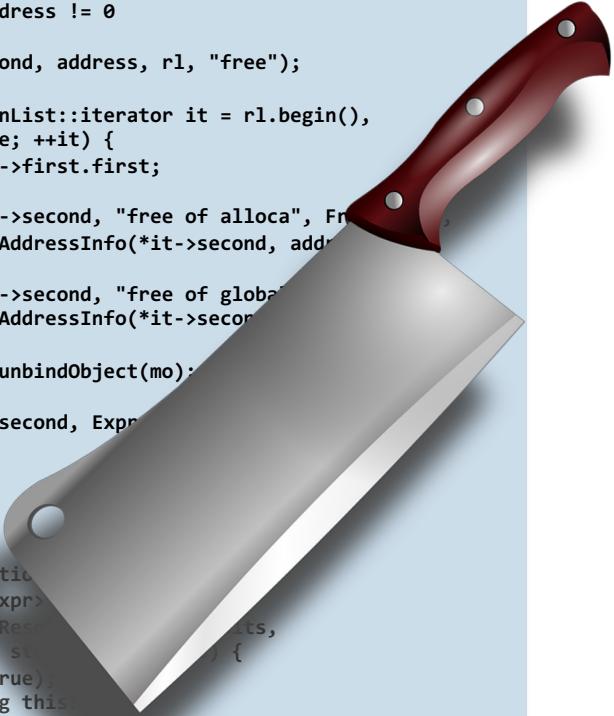
    for (Executor::ExactResolutionList::iterator it = rl.begin(),
        ie = rl.end(); it != ie; ++it) {
        const MemoryObject *mo = it->first.first;
        if (mo->isLocal) {
            terminateStateOnError(*it->second, "free of alloca", FnType::kUnknown);
            getAddressInfo(*it->second, address);
        } else if (mo->isGlobal) {
            terminateStateOnError(*it->second, "free of global", FnType::kUnknown);
            getAddressInfo(*it->second, address);
        } else {
            it->second->addressSpace.unbindObject(mo);
            if (target)
                bindLocal(target, *it->second, Expr::createPointer(0));
        }
    }
}

void Executor::resolveExact(Executor &executor,
                           ref<Expr> p,
                           ExactResolutionList &rl,
                           const std::vector<std::pair<FnType, std::string>> &fnTypes) {
    p = optimizer.optimizeExpr(p, true);
    // XXX we may want to be capping this
    ResolutionList rl;
    state.addressSpace.resolve(state, solver, p, rl);

    ExecutionState *unbound = &state;
    for (ResolutionList::iterator it = rl.begin(), ie = rl.end();
         it != ie; ++it) {
        ref<Expr> inBounds = EqExpr::create(p, it->first->getBaseExpr());
        StatePair branches = fork(*unbound, inBounds, true);

        if (branches.first)
            results.push_back(std::make_pair(*it, branches.first));

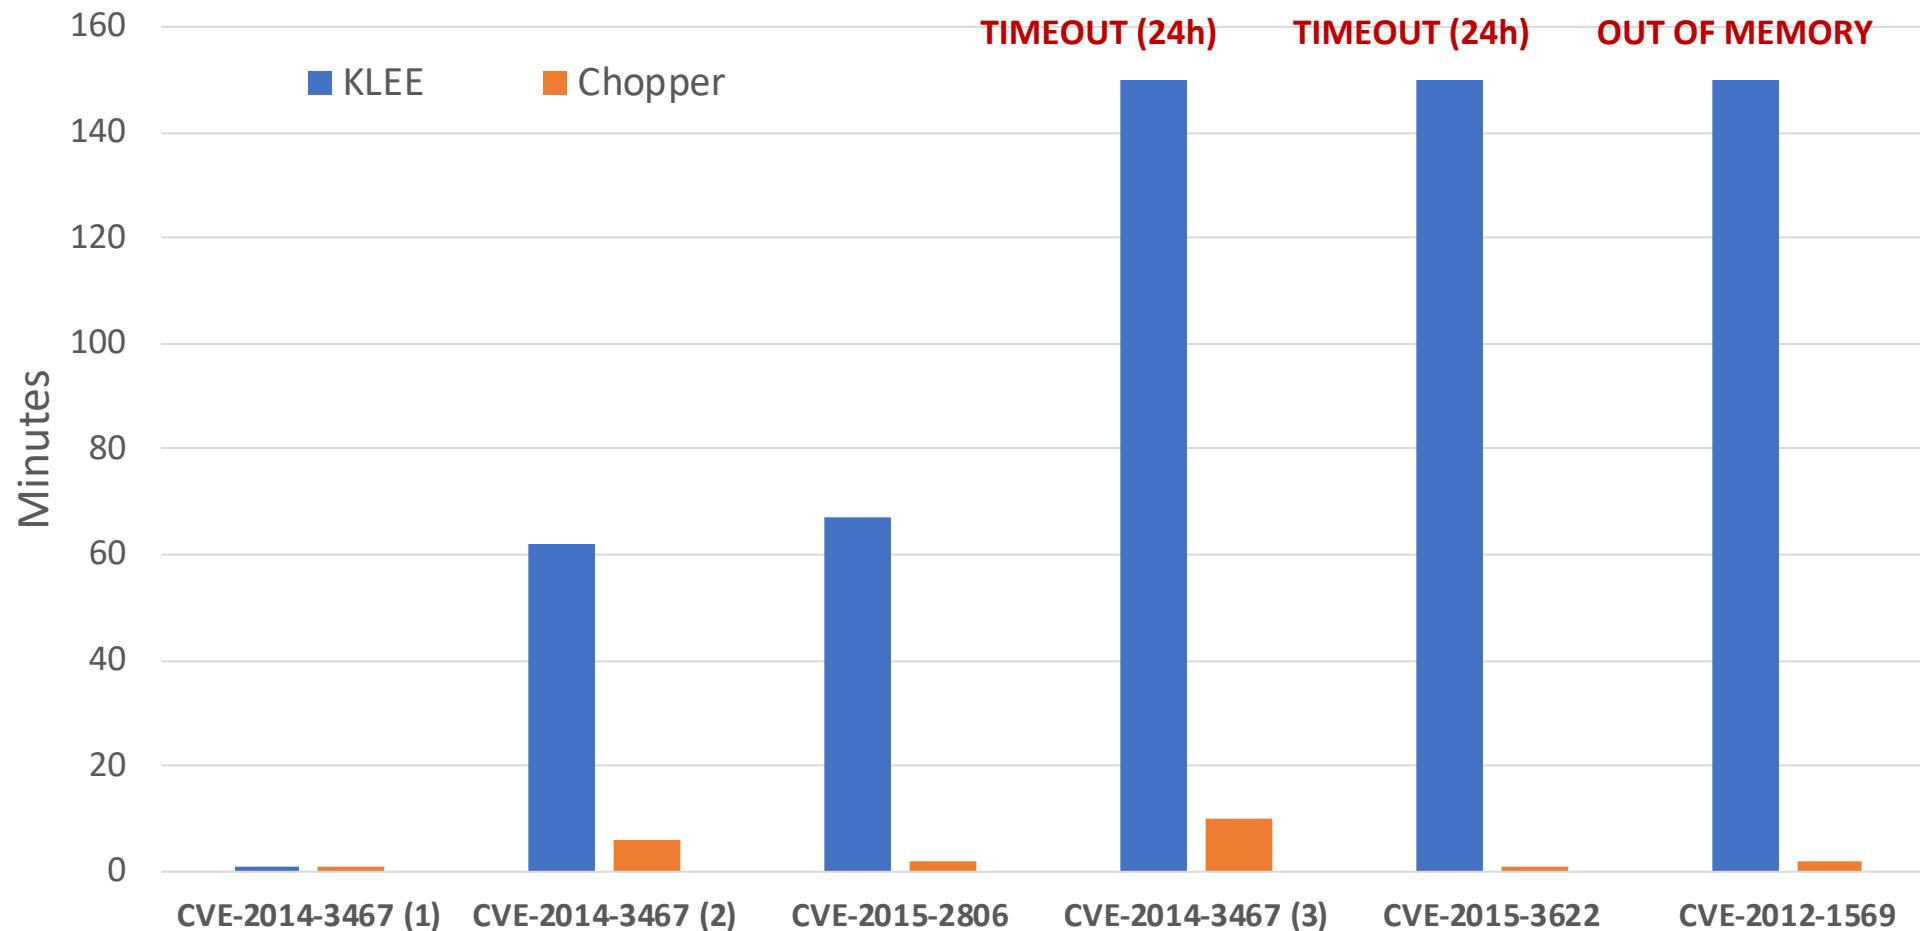
        unbound = branches.second;
        if (!unbound) // Fork failure
            break;
    }
}
```



Reproducing Security Vulnerabilities

- Benchmark: GNU libtasn1
 - ASN.1 protocol used in many networking and cryptographic applications, such as for public key certificates and e-mail
 - Considered 4 CVE security vulnerabilities, with a total of 6 vulnerable locations (out-of-bounds accesses)
- Goal:
 - Starting from the CVE report, generate inputs that trigger out-of-bounds accesses at the vulnerable locations
- Methodology:
 - Manually identified the irrelevant functions to skip
 - Time limit 24 hours, memory limit 4 GB

Reproducing Security Vulnerabilities



Effectiveness of Chopped Symbolic Execution

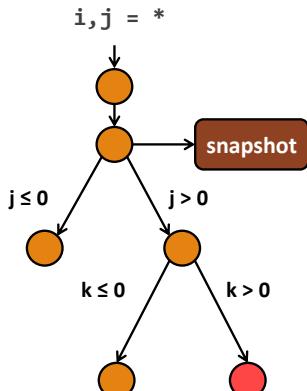
- Choice of code to skip
 - Application-specific, still work in progress
 - Some scenarios are easier to automate than others
 - Can always make different guesses and try them in parallel
- Precision of points-to analysis
 - Currently we use a flow-insensitive, context-insensitive and field-sensitive points-to analysis
 - Currently a single points-to analysis, in the beginning

From Whole-Program Analysis... ...To More Localized Tasks

- Most work on modern symbolic execution:
 - Whole-program test generation
 - Whole-program bug-finding
- More recently attention shifted to more localized tasks:
 - Patch testing
 - Debugging
 - Bug reproduction
 - Program repair
 - etc.
 - Which one is easier?
 - Opportunities of more localized tasks:
 - Prune a large part of the search space
 - Perform incremental reasoning/analysis
 - Use previous/correct version as an oracle
 - etc.

6

Chopped Symbolic Execution



```

void main() {
    f();
    if (j > 0) {
        if (y)
            target1;
    } else
        target2;
}
  
```

28

Experiments

SECURITY VULNERABILITY REPRODUCTION

```

address = optimizer.optimizeExpr(address, true);
StatePair zeroPointer = fork(state, Expr::createIsZero(address), true);
if (zeroPointer.first) {
    if (target)
        bindLocal(target, *zeroPointer.first, Expr::createPointer(0));
}
if (zeroPointer.second) { // address != 0
    ExactResolutionList r1;
    resolveExact(*zeroPointer.second, address, r1, "free");

    for (Executor::ExactResolutionList::iterator it = r1.begin(),
        ie = r1.end(); it != ie; ++it) {
        const MemoryObject *mo = it->first;
        if (!mo->isLocal) {
            terminateStateOnError(*it->second, "free of alloc");
            getAdddressInfo(*it->second, address);
        } else if (mo->isGlobal) {
            terminateStateOnError(*it->second, "free of global");
            getAdddressInfo(*it->second, address);
        } else {
            it->second->addressSpace.unbindObject(mo);
            if (target)
                bindLocal(target, *it->second, Expr::createPointer(0));
        }
    }
}

void Executor::resolveExact(Executor::RefExpr p,
                           ExactResolutionList &r1,
                           const std::vector<MemoryObject *> &branches,
                           p = optimizer.optimizeExpr(p, true);
// XXX we may want to be capping this
ResolutionList r1;
state.addressSpace.resolve(state, solver, p, r1);

ExecutionState *unbound = &state;
for (ResolutionList::iterator it = r1.begin(), ie = r1.end();
     it != ie; ++it) {
    RefExpr inBounds = EqExpr::create(p, it->first->getBaseExpr());
    StatePair branches = fork(*unbound, inBounds, true);
    if (branches.first)
        results.push_back(std::make_pair(*it, branches.first));
    unbound = branches.second;
    if (!unbound) // Fork failure
        break;
}
  
```

TEST SUITE AUGMENTATION

PATCH TESTING

Chopped Symbolic Execution [ICSE 2018]

David Trabish
Tel Aviv University
Israel
davivtra@post.tau.ac.il

Andrea Mattavelli
Imperial College London
United Kingdom
amattave@imperial.ac.uk

Noam Rinetzky
Tel Aviv University
Israel
maon@cs.tau.ac.il

Cristian Cadar
Imperial College London
United Kingdom
c.cadar@imperial.ac.uk

ABSTRACT

Symbolic execution is a powerful program analysis technique that systematically explores multiple program paths. However, despite important technical advances, symbolic execution often struggles to reach deep parts of the code due to the well-known path explosion problem and constraint solving limitations.

In this paper, we propose *chopped symbolic execution*, a novel form of symbolic execution that allows users to specify uninteresting parts of the code to exclude during the analysis, thus only targeting the exploration to paths of importance. However, the excluded parts are not summarily ignored, as this may lead to both false positives and false negatives. Instead, they are executed lazily, when their effect may be observable by code under analysis. Chopped symbolic execution leverages various on-demand static analyses at runtime to automatically exclude code fragments

the code with symbolic values instead of concrete ones. Symbolic execution engines thus replace concrete program operations with ones that manipulate symbols, and add appropriate constraints on the symbolic values. In particular, whenever the symbolic executor reaches a branch condition that depends on the symbolic inputs, it determines the feasibility of both sides of the branch, and creates two new independent *symbolic states* which are added to a worklist to follow each feasible side separately. This process, referred to as *forking*, refines the conditions on the symbolic values by adding appropriate constraints on each path according to the conditions on the branch. Test cases are generated by finding concrete values for the symbolic inputs that satisfy the *path conditions*. To both determine the feasibility of path conditions and generate concrete solutions that satisfies them, symbolic execution engines employ *satisfiability-modulo theory* (SMT) constraint solvers [19].