

# Practical Issues of Software Verification

Dirk Beyer



# State of the Art

- Much progress in MC theory & algorithms
- Practical issues in industrial applications
- Problems:
  - Large size of individual verification tasks
  - Large number of verification tasks

# Ideas

- Combine verification tools
- Reuse partial and intermediate results
- Witnesses for results validation
- Tests from Verification

# Classic Verification

C program

```
int main() {  
    int a = foo();  
    int b = bar(a);  
  
    assert(a == b);  
}
```

Specification



Verification  
Tool



**SAFE**

i.e., assertions  
cannot be violated



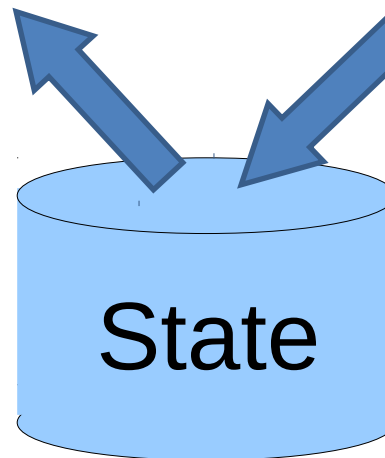


**UNSAFE**

# Stateful Verification

C program

```
int main() {  
    int a = foo();  
    int b = bar(a);  
  
    assert(a == b);  
}
```

Specification




**SAFE**

i.e., assertions cannot be violated

**UNSAFE**

# Applications of Stateful Verification

- Better performance by remembering successful (intermediate) results
- Regression Verification
- Certify results (verification witnesses)



# LMU Verification Center

# FSE 2012

## Conditional Model Checking: A Technique to Pass Information between Verifiers <sup>\*†</sup>

Dirk Beyer  
University of Passau  
Germany

Thomas A. Henzinger  
IST Austria  
Austria

M. Erkan Keremoglu  
Simon Fraser University  
Canada

Philipp Wendler  
University of Passau  
Germany

### ABSTRACT

Software model checking, as an undecidable problem, has three possible outcomes: (1) the program satisfies the specification, (2) the program does not satisfy the specification, and (3) the model checker fails. The third outcome usually manifests itself in a space-out, time-out, or one component of the verification tool giving up; in all of these failing cases, significant computation is performed by the verification tool before the failure, but no result is reported. We propose to reformulate the model-checking problem as follows, in order to have the verification tool report a summary of the performed work even in case of failure: given a program and a specification, the model checker returns a condition  $\Psi$ —usually a state predicate—such that the program satisfies the specification under the condition  $\Psi$ —that is, as long as the program does not leave the states in which  $\Psi$  is satisfied. In our experiments, we investigated as one major application of conditional model checking the sequential combination of model checkers with information passing. We give the condition that one model checker produces, as input to a second

### 1. INTRODUCTION

Model checking is an automatic search-based procedure that exhaustively verifies whether a given model (e.g., labeled transition system) satisfies a given specification (e.g., temporal-logic formula) [12][33]. Since model checking of software is an undecidable problem, there are three possible outcomes of the analysis process: (1) the program satisfies the specification, (2) the program does not satisfy the specification, and (3) the model checker fails. The first outcome can be obtained by the model checker if the abstract model that was computed for the program is sufficient to prove the program correct under the given specification. This outcome can be accompanied by a proof certificate [23]. The second outcome can be obtained by the model checker if an abstract counterexample path is found and can be proven feasible, i.e., a bug that can actually occur in the program. This outcome is usually accompanied by the violating program part in the form of program source code, and sometimes test input to reproduce the error at run-time [4]. The third outcome usually occurs if the model checker runs out



# Software Verification

C program

```
int main() {  
    int a = foo();  
    int b = bar(a);  
  
    assert(a == b);  
}
```

Specification



Verification  
Tool



**SAFE**

i.e., assertions  
cannot be violated




**UNSAFE**

# Problem:

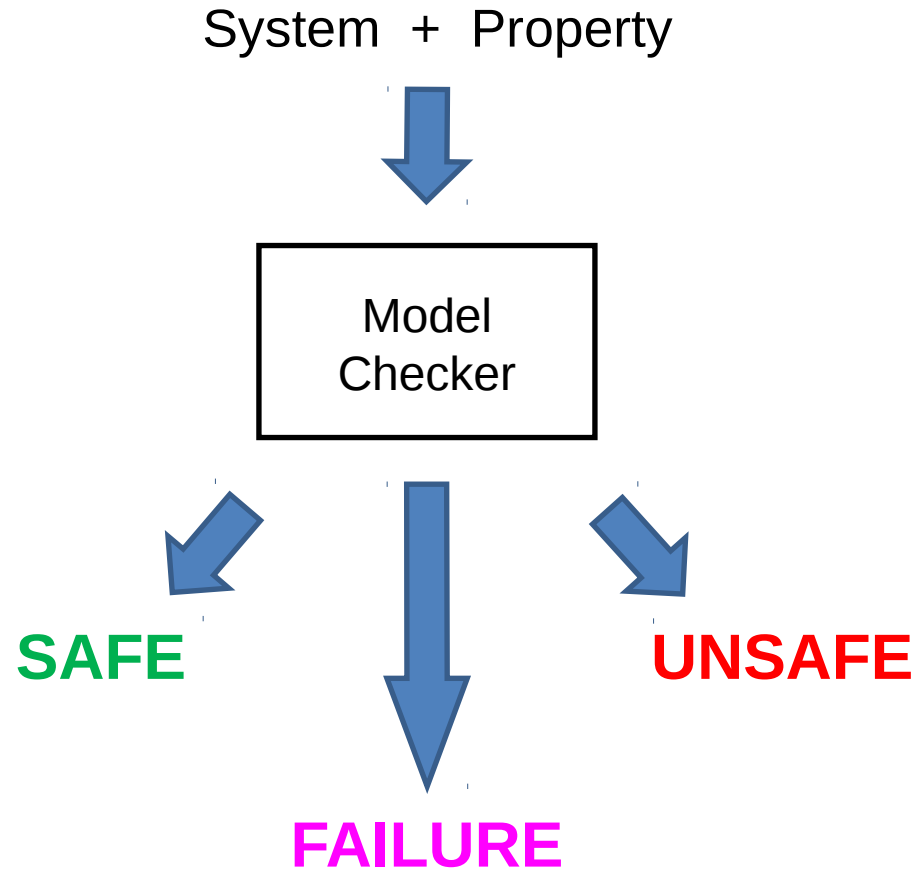
## Single Analysis not Effective

```
1 void main() {
2     if (nondet_int()) {
3         int i;
4         for (i = nondet_int(); i < 1000000; i++) {
5             // ...
6         }
7         assert(i >= 1000000);
8
9     } else {
10        int x = 5;
11        int y = 6;
12        int r = x * y;
13        assert(r >= x);
14    }
15 }
```


# Model Checking



# Classic Model Checking




# Classic Model Checking



# Classic Model Checking

System + Property



**FAILURE**

- Timeout
- Out of memory
- Crash of component
- Operand exception

Enormous amounts of resources **wasted!**

# Conditional Model Checking




FSE 2012, joint work with Tom Henzinger,  
Erkan Keremoglu, Philipp Wendler



# Conditional Model Checking

System + Property



$\Psi$

**(“SAFE under Condition  $\Psi$ ”)**

- Examples:
- $\Psi = \text{true}$ : previous SAFE
  - $\Psi = \text{false}$ : previous UNSAFE
  - general: condition for safety




# Conditional Model Checking

System + Property

Condition  $\Psi_0$

Directs the analysis to parts to analyze




$\Psi$

**(“SAFE under Condition  $\Psi$ ”)**

- Examples:
- $\Psi = \text{true}$ : previous SAFE
  - $\Psi = \text{false}$ : previous UNSAFE
  - general: condition for safety

# Applications of Conditional Model Checking



# Back to Our Example

```
1 void main() {
2     if (nondet_int()) {
3         int i;
4         for (i = nondet_int(); i < 1000000; i++) {
5             // ...
6         }
7         assert(i >= 1000000);
8
9     } else {
10        int x = 5;
11        int y = 6;
12        int r = x * y;
13        assert(r >= x);
14    }
15 }
```

# Back to Our Example

To show:

$$M \models \Phi$$

In this case:

$$\Phi = \Phi_1 \ \& \ \Phi_2$$

with  $\Phi_1 =$  “loop is correct”

and  $\Phi_2 =$  “multiplication is correct”

# Idea: Decompose!


- Verify  $\Phi_1$  (“loop is correct”)
  - use predicate analysis
- Verify  $\Phi_2$  (“multiplication is correct”)
  - use explicit-state analysis
- Final result:  $\Phi$  verified

# Using CMC with Input Conditions

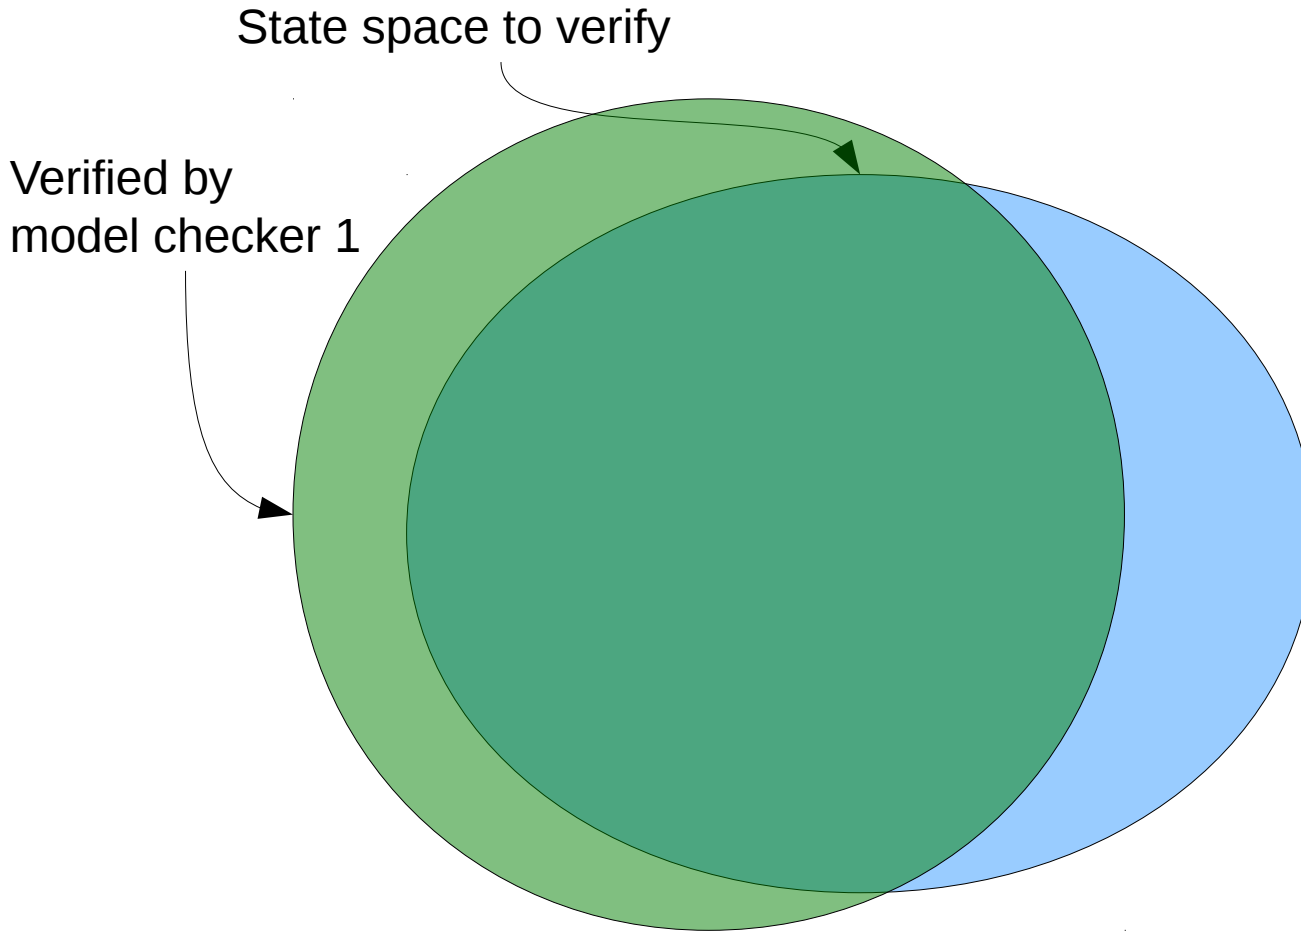
- Tell model checker what to verify
- In our example:
  - For conditional model checker 1: verify  $\Phi_1$
  - For conditional model checker 2: verify  $\Phi_2$
  - Full verification possible

# More General:

State space to verify




# More General:





# More General:



# Further Input Conditions

- Limit resources
  - Time
  - Memory
  - Model Checker will not crash, but terminate itself and give useful result
- Restrict the search
  - Loop bounds (a.k.a. “bounded model checking”)
  - Path length
  - Time spent on path
  - ...

# Output Conditions

- Dump partial result if analysis didn't finish
  - Output cond. summarizes what could be verified
- Explicitly state assumptions used by MC
  - Example: “variable **x** does not overflow”
- Purpose:
  - Give information to the user
  - Verify condition with other methods (testing, manual proofs, ...)
  - Comparison of checkers (weaker output condition is better)



# Sequential Composition

- In our example,  
we told the model checkers what to verify
- Now let them find out automatically!
- Conditional model checker 1 verifies  
what it can verify
- Conditional model checker 2 verifies  
remaining parts

# Sequential Composition



- Use input condition to limit resource usage of first analysis
- Use output condition as input condition for next model checker
- Iterate until finished (or run out of tools)

# Sequential Composition




State space to verify

# Sequential Composition




State space  
partially verified

# Sequential Composition





# Sequential Composition



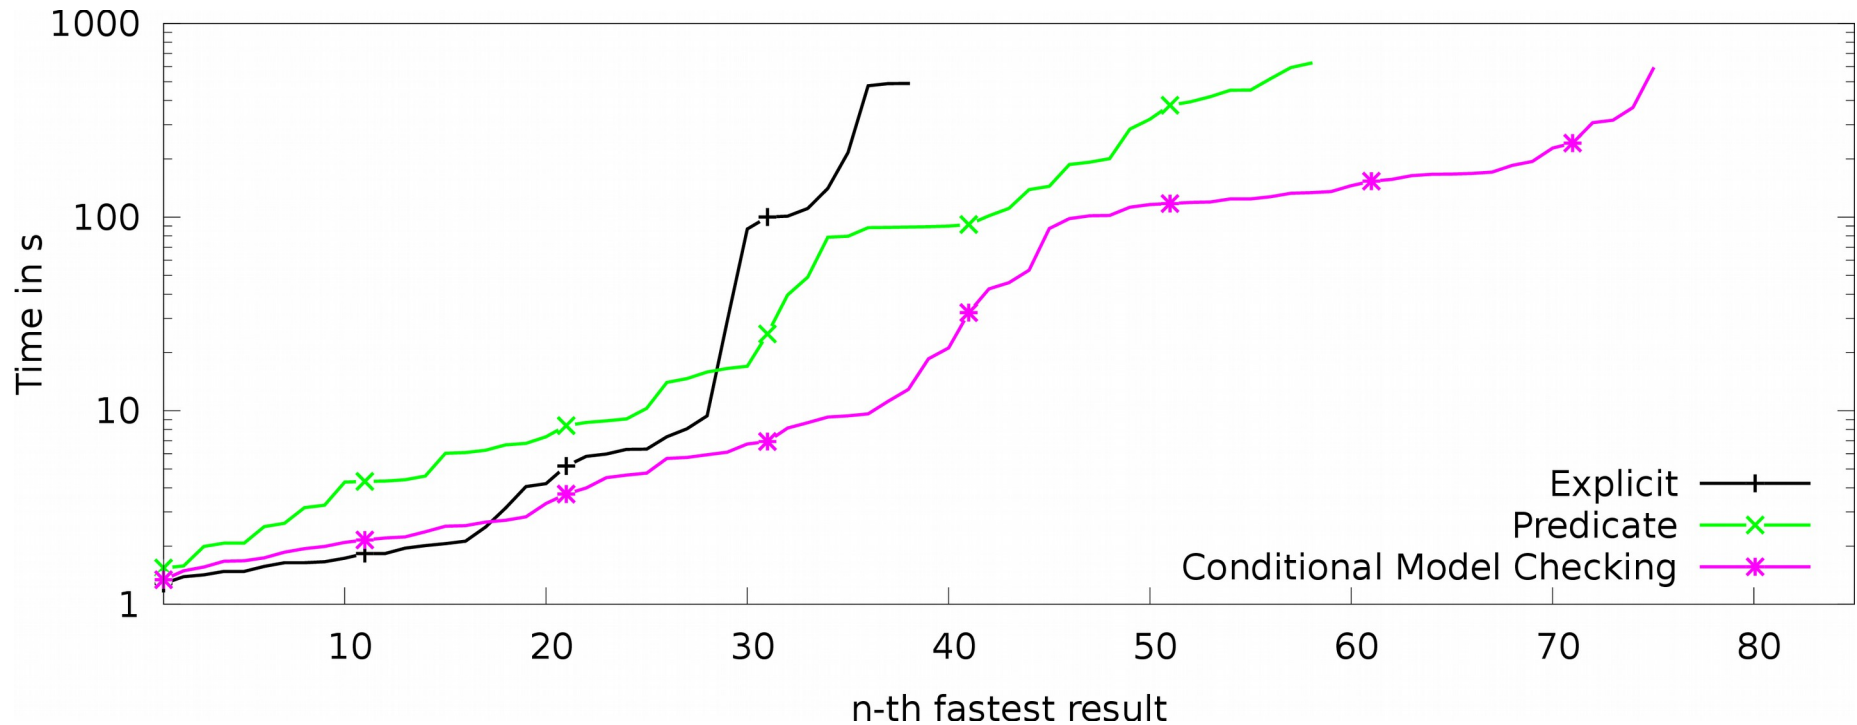
# Experiment: Sequential Composition

- Implemented Conditional Model Checking in CPAchecker
- 85 C programs based on “hard” programs of Software Verification Competition 2012
- 15 min time, 15 GB RAM

# Experiment: Sequential Composition

- A: Explicit-value analysis
- B: Predicate analysis
- C: Conditional model checking
  - First: explicit-value analysis  
with input condition: time limit = 100s
  - Second: predicate analysis  
with output condition of first analysis  
as input condition

# Experiment: Sequential Composition




→ Sequential composition  
solves more problems and is faster

# Experiment:

## Sequential Composition

- A: Explicit-value analysis ; predicate analysis
- B: Explicit-value analysis ; predicate analysis
  - Input condition for first analysis:  
time limit = 100s
- C: Conditional model checking
  - First: explicit-value analysis  
with input condition: time limit = 100s
  - Second: predicate analysis  
with output condition of first analysis  
as input condition

# Experiment: Sequential Composition



→ Using conditional model checking for sequential composition is better

# Summary Part 1

## Conditional Model Checking:




- Terminates with useful results  
(no crashes)
- Enables partial / compositional verification
- Effective sequential composition  
(solve harder problems)
- Unified view on existing approaches

# Stateful Verification

C program

```
int main() {  
    int a = foo();  
    int b = bar(a);  
  
    assert(a == b);  
}
```

Specification



**SAFE**

i.e., assertions cannot be violated

**UNSAFE**



# Towards Reusing Information

- Context: CEGAR-based verification
- Abstract model has to be constructed every time a verification task is started
  - Refinements of Precision
  - Reconstruction and Pruning of ARG

# Precision Reuse for Efficient Regression Verification

(Published in Proc. ESEC/FSE 2013, ACM.)

Dirk  
Beyer



Stefan  
Löwe




Evgeny  
Novikov



Andreas  
Stahlbauer



Philipp  
Wendler





Linux Driver Verification

# Example

Revision	Commit Message	Safe?
3	Implement button detection support	✗
4	Free MICDET IRQ on error during probe	✗
5	fix typos in extcon-arizona	✗
6	Use bypass mode for MICVDD	✗
7	Merge tag 'driver-core-3.6' of ...	✗
8	unlock mutex on error path in ...	✓
9	remove use of devexit	✗
10	remove use of devinit	✗
11	remove use of devexit p	✗
12	Merge tag 'pull req 20121122' of ...	✓

# High Resource Consumption!

Software Verification is expensive


Verifying all **safety properties** for all **revisions** of a software ...

500 drivers  
\* 60 properties  
\* 2 before/after  
= 60 000 verification tasks  
\* 10 seconds/verification task  
= 600 000 seconds

≈ 7 days



... is really expensive



# LMU Verification Center

# Reuse of Verification Results

## Drawbacks of existing approaches

- **Too large**: space on disk, time for loading
- **Too sensitive** to changes between revisions
- **Too complex**: modification of the verification algorithm



➔ Reuse the “precision”

# Precision $\pi$

Defines the **level of abstraction** within an abstract domain:

**Information** that an abstraction-based analysis has **to track** to prove a property.



# Examples for Precision

- Predicate Analysis

$$\pi = \{a > 0, k == 1 \wedge e == 0\}$$

**Set of predicates** used to compute boolean abstractions

- Explicit-Value Analysis

$$\pi = \{a, k, e\}$$

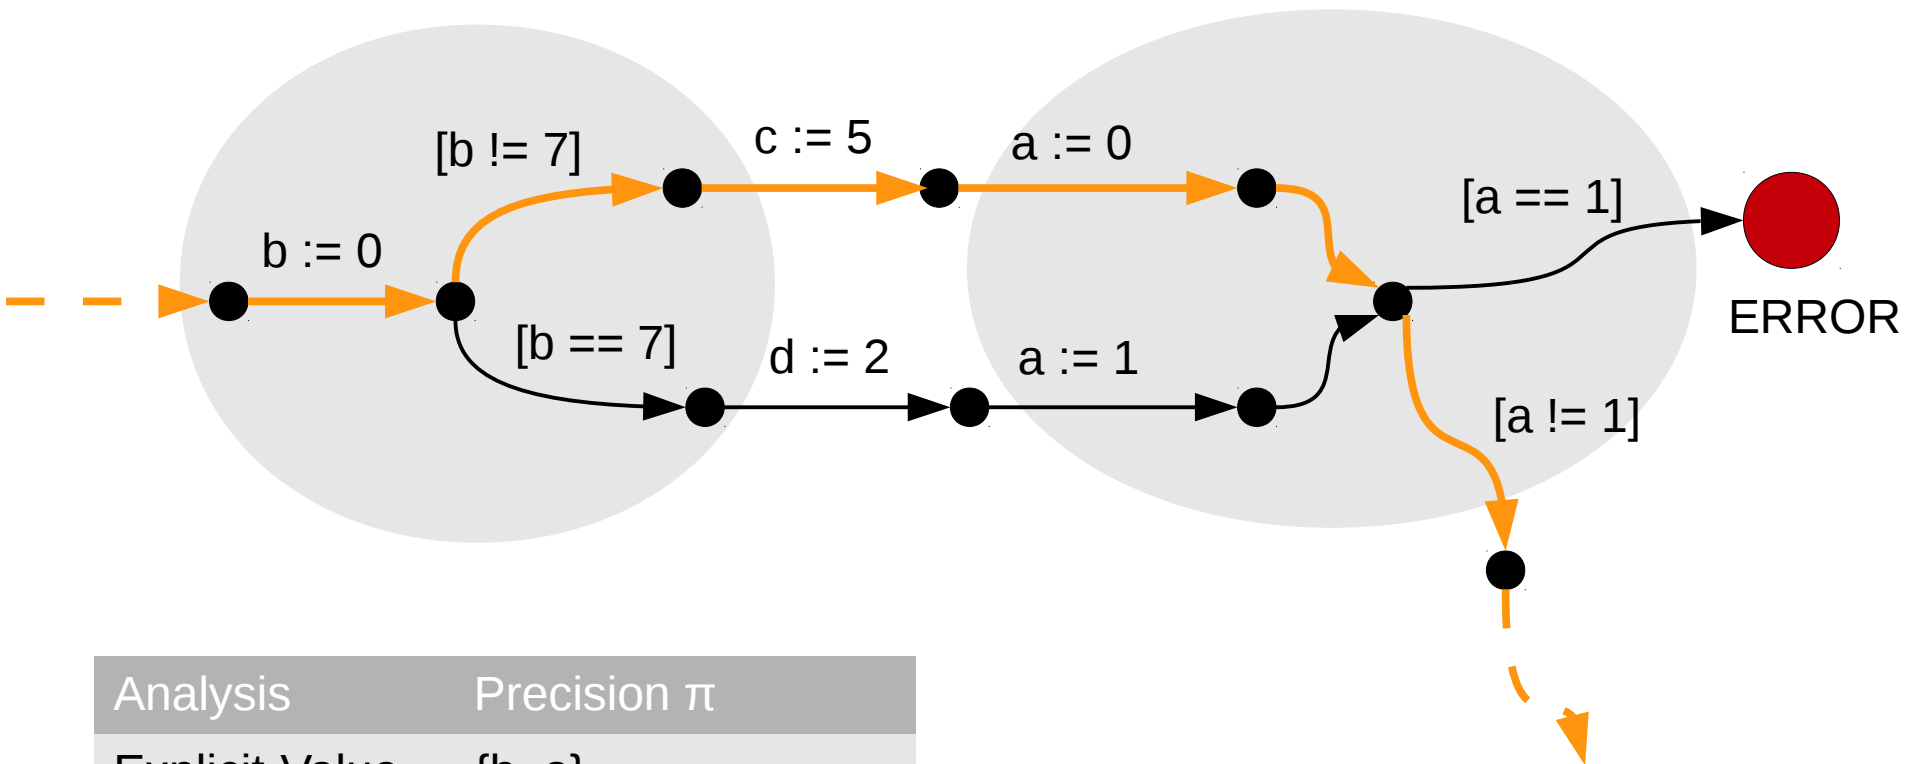
**Set of variables** for which the explicit value has to be tracked

- Shape Analysis

$$\pi = \{p1, p2\}$$

**Set of pointer variables** to track


# Example




Analysis	Precision $\pi$
Explicit-Value	{b, a}
Predicate	{b == 7, a == 1}

# Advantages of Reusing Precisions

- ✓ **No modification** of the verification algorithm
- ✓ **Easy to extract** from model checkers
- ✓ **Small** memory footprint
- ✓ **Low sensitivity** to changes in the input programs




# CEGAR



$$\pi_{i+1} = \pi_i \cup \text{Interpolants}_{i+1}$$

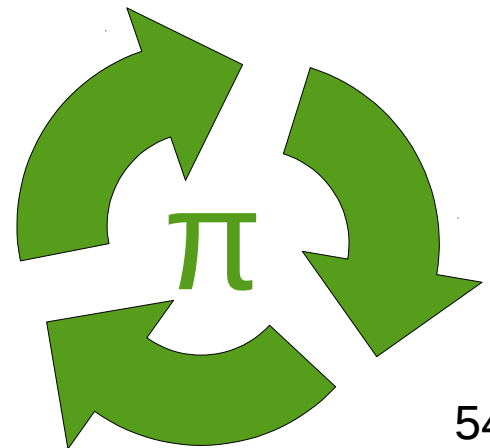
# CEGAR + Reuse



$$\pi_{i+1} = \pi_i \cup \text{Interpolants}_{i+1}$$

# Advantages of Reusing Precisions

- ✓ **No modification** of the verification algorithm
- ✓ **Easy to extract** from model checkers
- ✓ **Small** memory footprint
- ✓ **Low sensitivity** to changes in the input programs



# Implementation

<http://cpachecker.sosy-lab.org>

- Implemented in CPAchecker

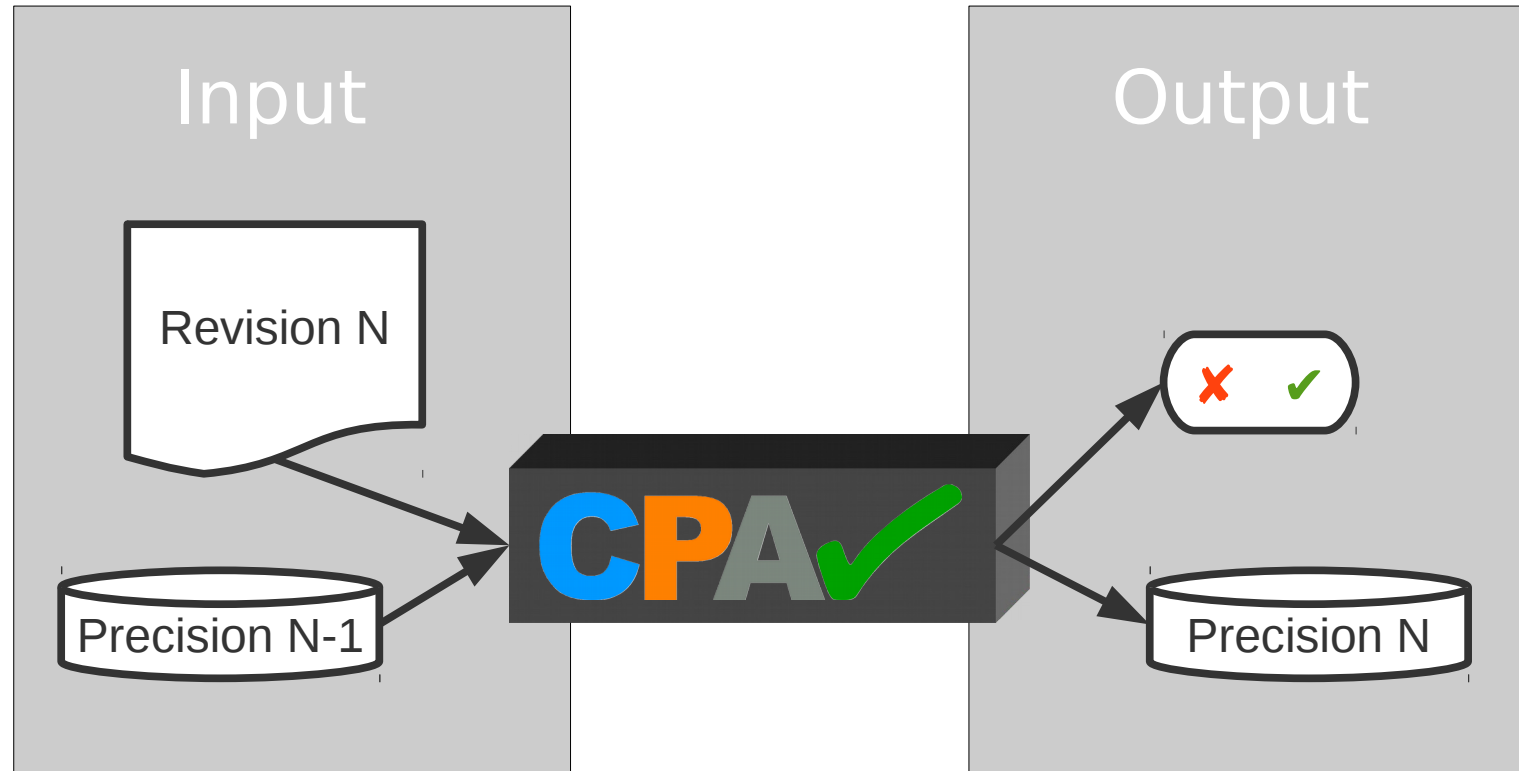
- Predicate Analysis
- Explicit-State Analysis



- Common to both analyses:

- Lazy abstraction
- CEGAR
- Construct an abstract reachability graph

# Workflow






# Stateful Verification

C program

```
int main() {  
    int a = foo();  
    int b = bar(a);  
  
    assert(a == b);  
}
```

Specification



Verification Tool




**SAFE**

i.e., assertions cannot be violated



**UNSAFE**



State

# Storing Precisions

## Explicit-State Analysis

```
* :  
lock  
  
main f:  
x
```

## Predicate Analysis

```
(declare-fun |lock|() Real)  
(declare-fun |x|() Real)  
(define-fun t1() Bool (= |lock| 0))  
(define-fun t2() Bool (<= |x| 1))  
  
* :  
(assert t1)  
  
main f:  
(assert t2)
```

Really simple! **Dump the precision**

# Benchmark Suite


- Derived from industrial code (Linux kernel)
  - 4193 verification problems
  - 62 Linux device drivers
  - 1119 revisions  
spanning more than 5 years of development
- Publicly available

<http://sosy-lab.org/~dbeyer/cpa-reuse/>

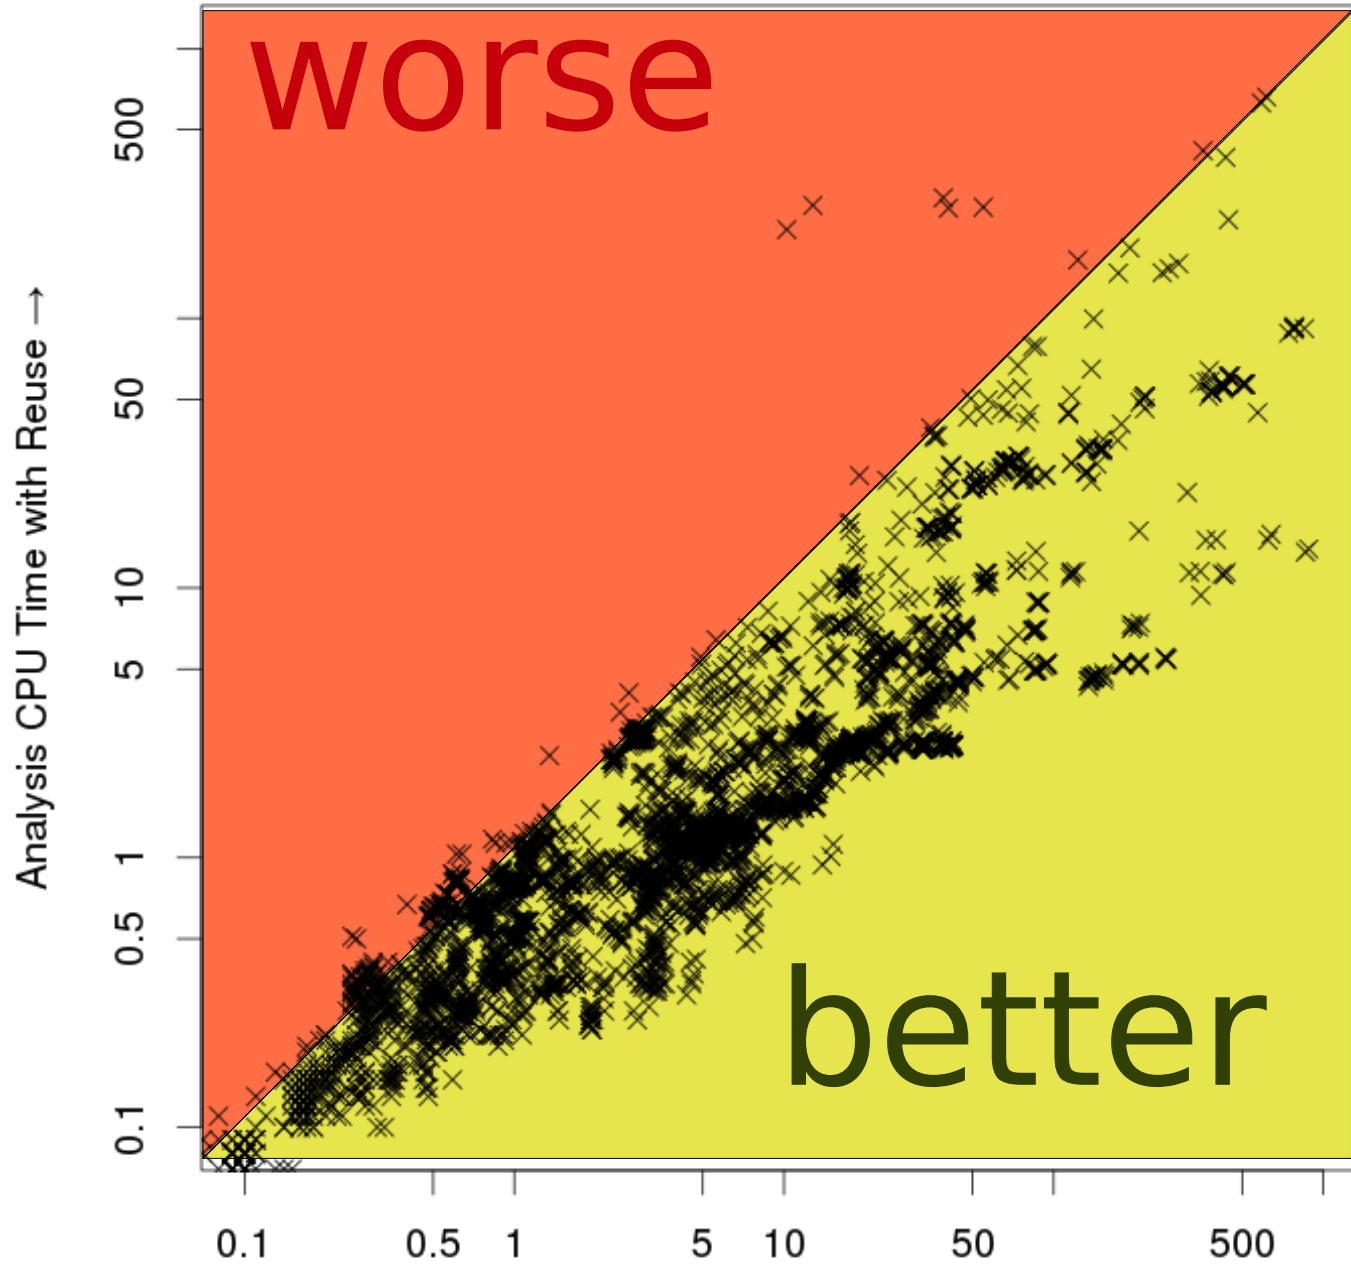
# Benchmark Setup

- Processor: Intel i7 3.4 GHz Quad Core
- Time limit: 15 minutes
- Memory limit: 15 GB


= Setup of the Intl. Competition on Software Verification



# Results for Predicate Analysis



# Results for Predicate Analysis



# Tasks	4 193
Analysis CPU Time without Reuse	83 000
Analysis CPU Time with Reuse	23 000
Speedup	<b>4.3</b>
Solved	4 048 + 30

# Summary – Part 2

Precision reuse has a significant positive effect!



- Drastically improves performance
  - Reduces the number of refinements
- More problems can be solved
- Low sensitivity to changes in the program code



# Reusing Witnesses

- Learn from previous proofs
- If you know a previous error path,
  - check this first, try to “re-play”
- If you know a previous proof,
  - try to “re-validate”, watch for changes