

From Finding Bugs to Certifying their Absence

Sagar Chaki

September 21, 2006

Motivation

We are increasingly **reliant** on **software** despite continuing **problems** controlling software **quality**.

Human and financial

- Therac-25 X-Ray machine
 - At least 6 known incidents with 5 fatalities
- Ariane 5 loss (~\$640 million)
- Pentium recall (~\$450 million)

Recurring **maintenance** costs – security patches

Consumer **confidence**

- Airport shutdown due to malfunction in baggage X-Ray

Development costs and **delayed** time to market

Finding Bugs

Experience points us to **common causes** of bugs

- Improper use of **shared resources**
- Failure to comply with **API policies**
- Poor **coding** practices
- Buffer **overflow**

Nevertheless, bugs remain **undetected** and with unknown effects **despite** expensive **quality assurance** practices

- Effects may not be observed, e.g., during testing
- Even if observed, difficult to **reproduce** and eradicate

Bug-Finding Technology

For generic problems – good, but only partial solutions

- Compilers
- Commercial code analysis tools

For specific problems

- Testing – incomplete coverage, focus on effect
- Code inspection – inconsistent results

We need a means to get **consistent** results with more complete **coverage** when looking for known causes

Software Model Checking

Key ideas:


- Exhaustive analysis
- Automated abstraction of conservative finite models
- Policy independent

Strengths:

- Usable at industrial scale*
- Packages tool complexity and deep domain expertise

Limitations:

- Scale will always be a factor
- Asymmetric confidence



* <http://www.microsoft.com/whdc/devtools/tools/sdv.mspix>

Model Checking Successes

Typically applied by **experts** for **specific products**

- Cache coherence protocols
 - IEEE Futurebus+
- Microprocessors
 - Intel, IBM, Motorola
- Telecommunications
 - Lucent PathStar switch

Microsoft Static Driver Verifier (SDV) is a notable exception and a good example of **big impact**, **broad application** of software model checking

Microsoft Static Driver Verifier

Goal: reduce “blue screen of death” occurrences.

Many failures traced to **improper interaction** by third-party **device drivers** with the Windows kernel.

Solution: **identify** correct interaction policies and **automate** violation detection with software model checking.

Restriction to known policies allowed **optimizations** for scale and **black box use** by any engineer.

SDV has been used as part of driver certification since 2003 and is now publicly available

- <http://www.microsoft.com/whdc/devtools/tools/sdv.msp>

Our Experience

Robotics line communications library

- Widely understood domain
- Able to identify meaningful policies
- Found several important bugs

Embedded Operating System (6000 LOC)

- Less expertise, more general policies
- Still successful at finding bugs

Focused application of model checking for specific issues yields practical improvements

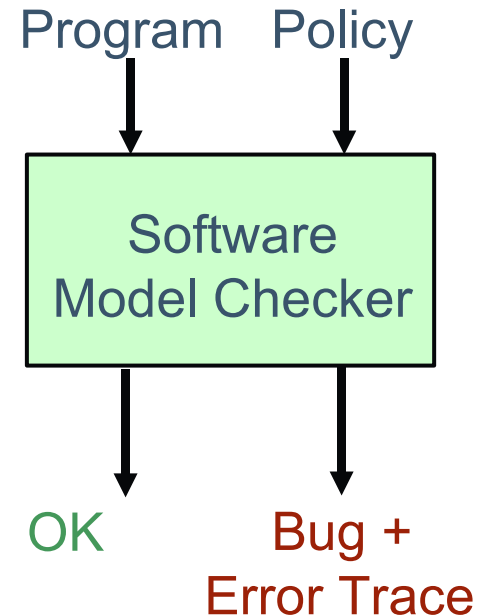
- **Buffer overflow** identification with fewer false positives

Verification and Trust

Software model **checking** alone is **inadequate** for **safety-critical** systems

- Checker itself a complex system
 - Policies pass because checker is buggy
- Difficult to detect this situation
 - Checker only says “OK”
 - Have to trust the checker

Goal: **Move** the model checker **out** of the Trusted Computing Base (TCB)



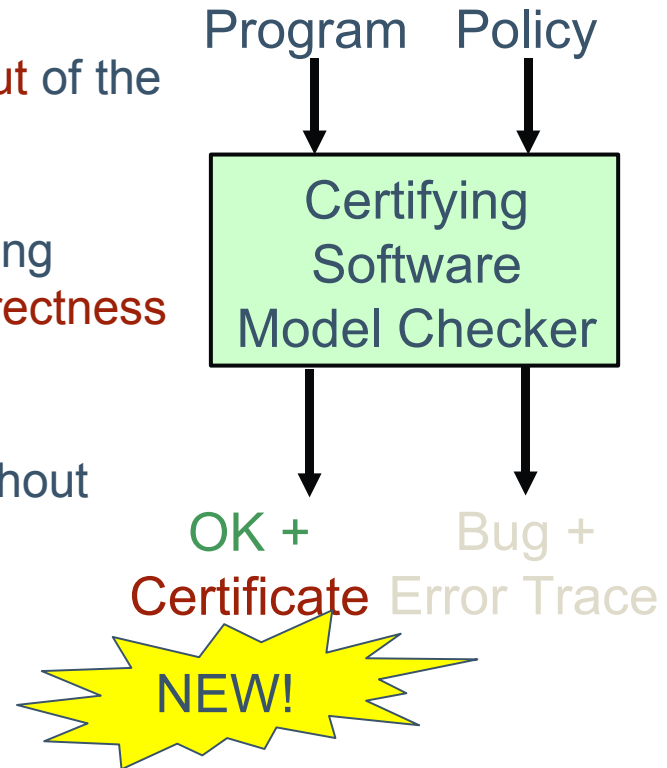
From Verification to Certification

Have to **trust** model checker if “OK”

Goal: **Move** the model checker **out** of the Trusted Computing Base (TCB)

Idea: Use **certifying** model checking

- Model checker generates a **correctness** certificate with “OK”
- Certificate can be **validated** **independently** and **objectively** without trusting the model checker
- **Symmetric** confidence



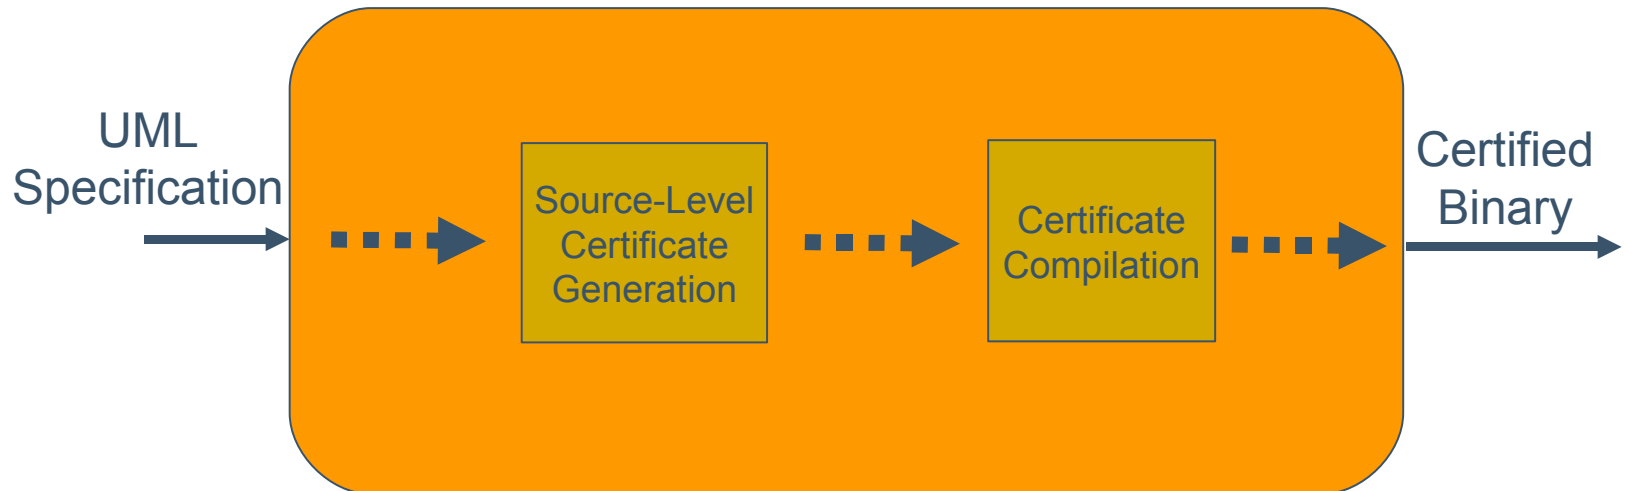
Certifying Binaries

Certifying software model checking certifies **source** code

- Still have to **trust** the **compiler** – another complex system

Can eliminate the compiler from the TCB

- By “**compiling**” the **certificate** generated by the model checker
- Computationally **inexpensive**



Rest of the Talk

A framework to generate certified assembly code from UML specifications in a completely automated manner

- Step 1: Certify source code
- Step 2: Push certification to assembly code
- No need to trust the mechanisms (humans or tools) involved in the certification process

Implementation and experimental validation

- Demonstration to follow later on in the presentation

Certifying a Program

The **program** is like a maze

- Can be certified by the following two artifacts

(1) Invariants

- Conditions that always hold during execution
- Similar to the **breadcrumbs**

(2) Proof of the following facts:

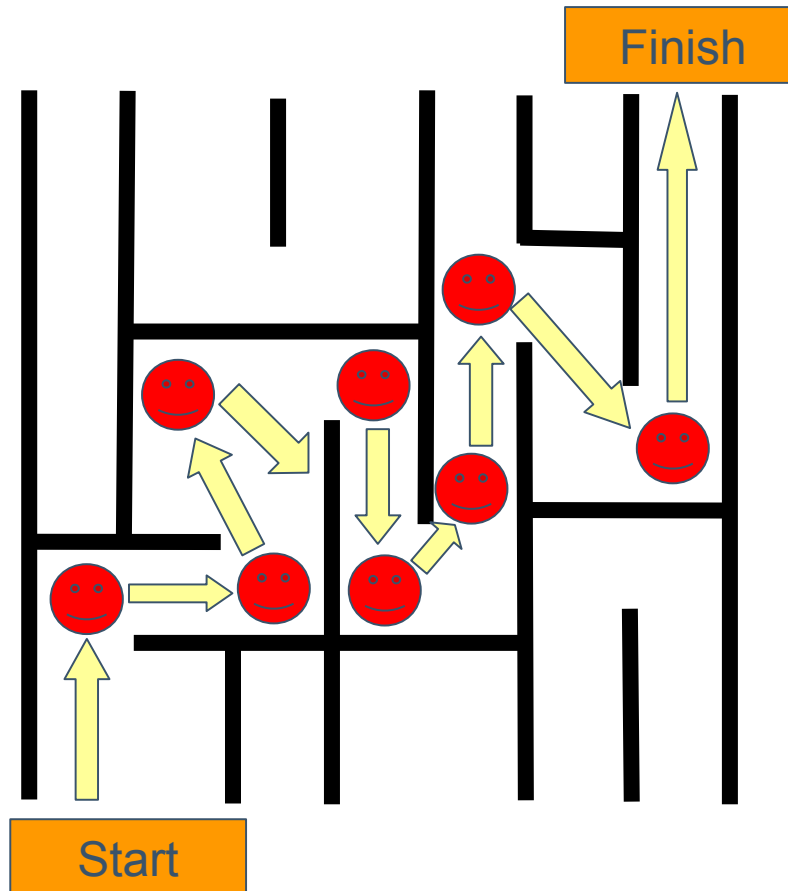
- Each invariant implies the next
- Invariants collectively imply the policy
- Similar to the **arrows**

Verification
Condition
(VC)

Invariants + Proof = Certificate

- Certificate can be validated objectively

Certificate Invalidated by Other Changes



Certificate
changed

Certification Protocol

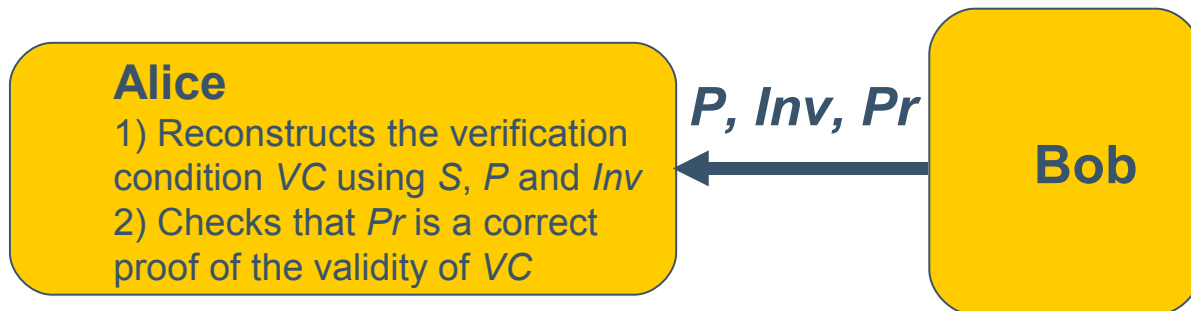
Bob wants to convince Alice that his program P satisfies a policy S

- Policy S may or may not be supplied by Alice

Bob supplies Certificate = Invariant (Inv) + Proof of VC (Pr) along with P

Alice **checks** that the certificate is correct

- Typically much easier than constructing the certificate



The VC construction procedure is public and deterministic. So Bob can construct the exact VC as Alice in order to produce a correct proof. Alice does not have to trust Bob.

Basic Definitions

Two types of policy

- **Safety** : something bad never happens
 - a channel is never used before being secured
- **Liveness** : something good eventually happens
 - a channel is eventually secured

Invariant : A condition that is always true whenever the program execution reaches a specific program location

- Used to certify safety

Ranking function : A mechanism for imposing an ordering between various program states

- Used to certify liveness

Example

```
#define CHANNEL_SECURED (n == 10)

void main()
{
    int n = 0;
    [n >= 0 && n <= 10]
    while(n < 10)
        [n >= 0 && n < 10]
        n = n + 1;
    [n == 10]
    assert(CHANNEL_SECURED);
    use_channel();
    return;
}
```

invariants

**liveness: this assertion
is eventually checked**

**safety: this assertion
never fails**

Example

```
#define CHANNEL_SECURED (n == 10)

void main()
{
    int n = 0;
    [n >= 0 && n <= 10]
    while(n < 10)
        [n >= 0 && n < 10]
        n = n + 1;
    [n == 10]
    assert(CHANNEL_SECURED);
    use_channel();
    return;
}
```

invariants

**safety policy below
can be certified
using invariants above**

**safety: this assertion
never fails**

Certification Protocol

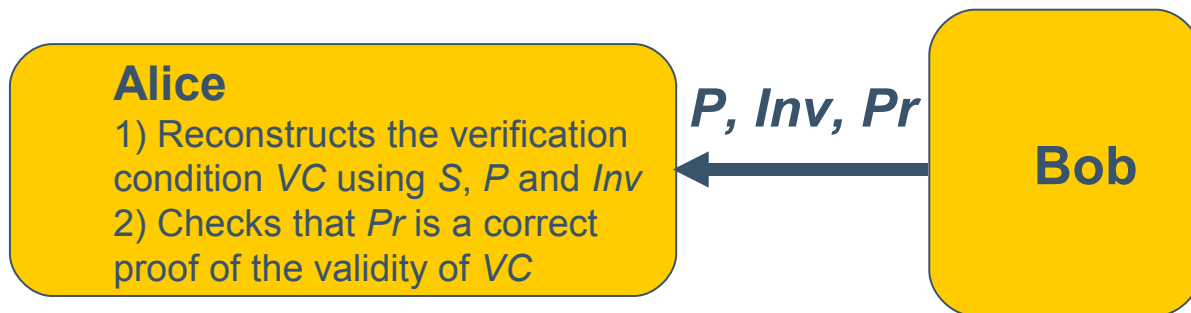
Bob wants to convince Alice that his program P satisfies a policy S

- Policy S may or may not be supplied by Alice

Bob supplies Certificate = Invariant (Inv) + Proof of VC (Pr) along with P

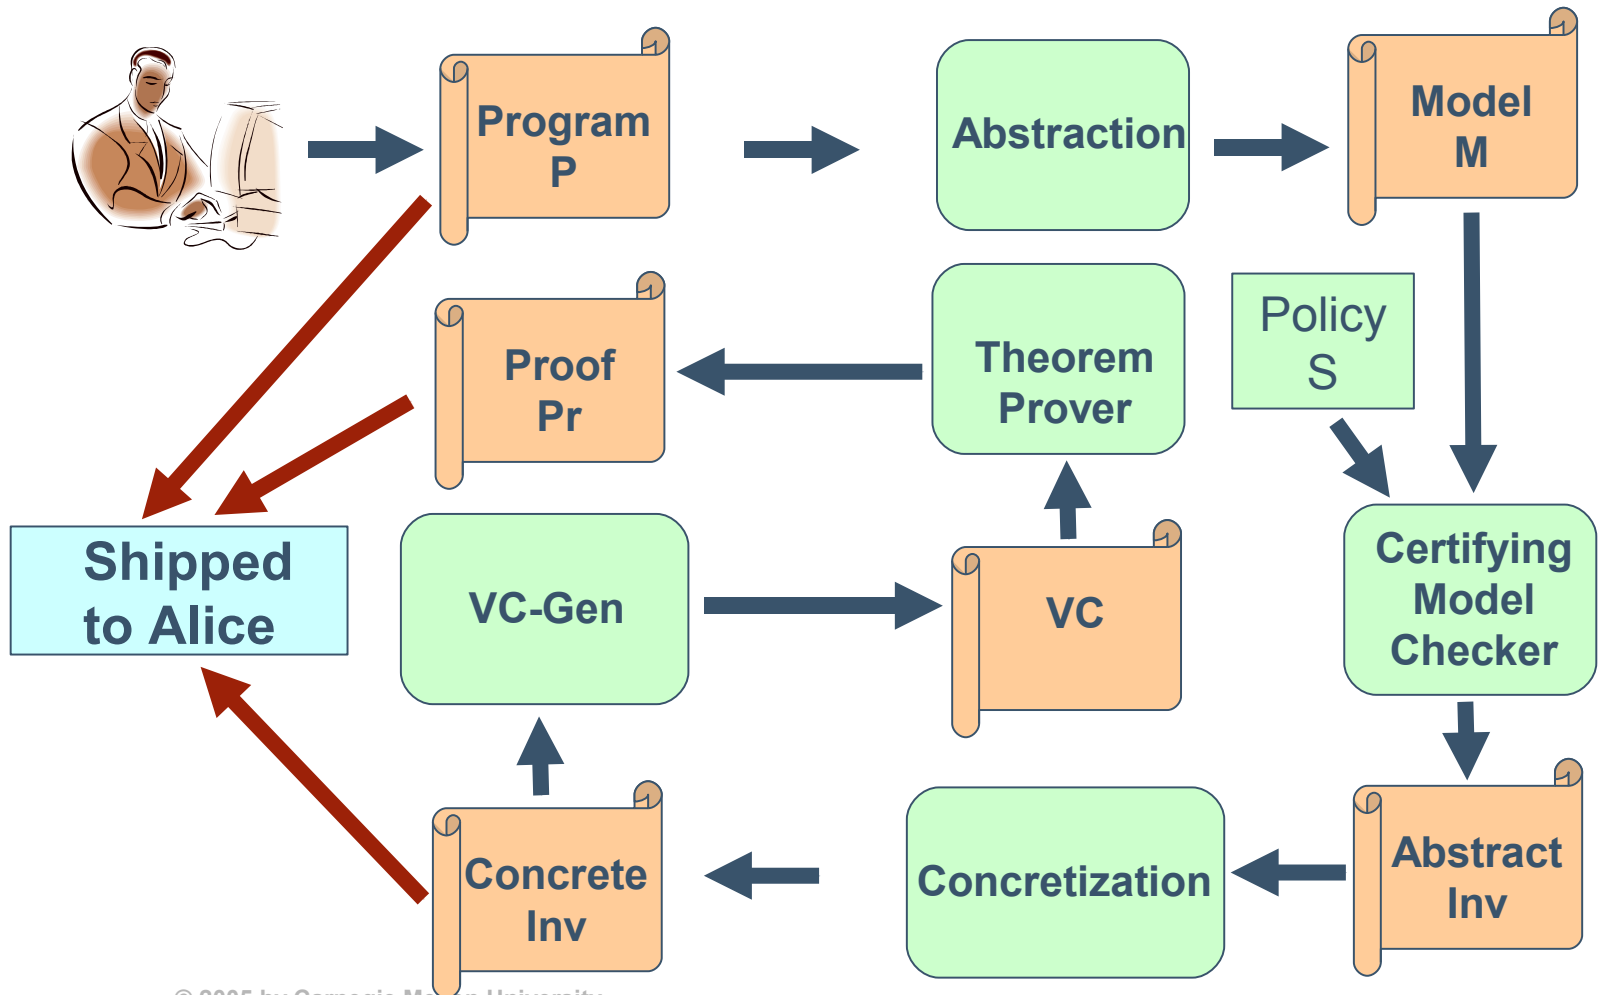
Alice **checks** that the certificate is correct

- Typically much easier than constructing the certificate



Let's look more closely at Bob

Bob: Internal Details



SAT-based Software Certification

Used Boolean satisfiability to generate compact proofs that certify safety and liveness policies of C programs

- were able to certify correct locking behavior in device drivers and an embedded OS
- proofs were 100 -- 100,000 times smaller

SAT-based Software Certification, S. Chaki,
Proceedings of Tools and Algorithms for the
Construction and Analysis of Systems (TACAS),
2006

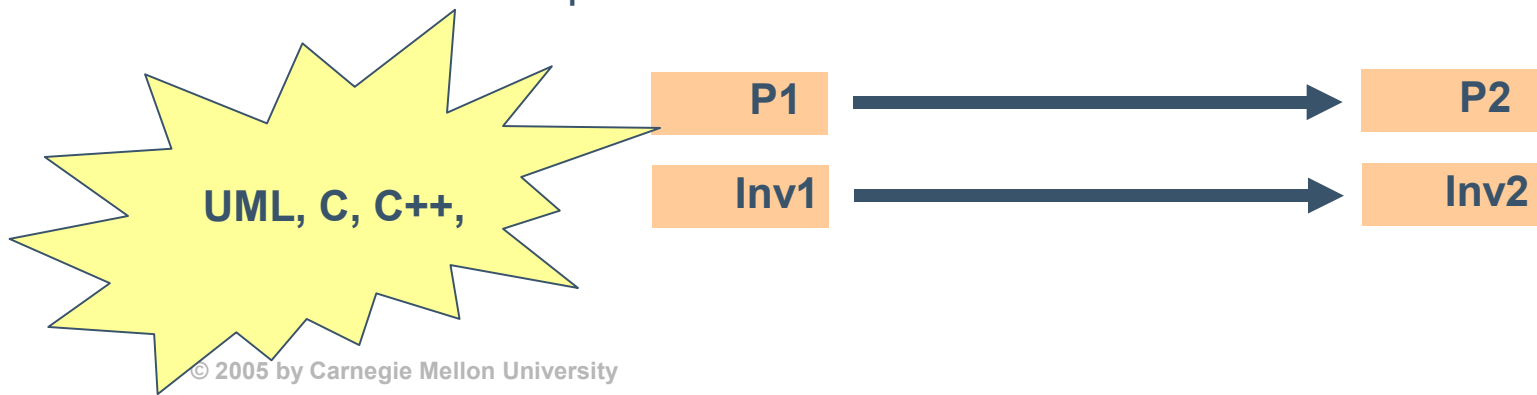
Certifying Binaries

A binary are just another kind of program. The difficulty is obtaining appropriate invariants from low-level instructions.

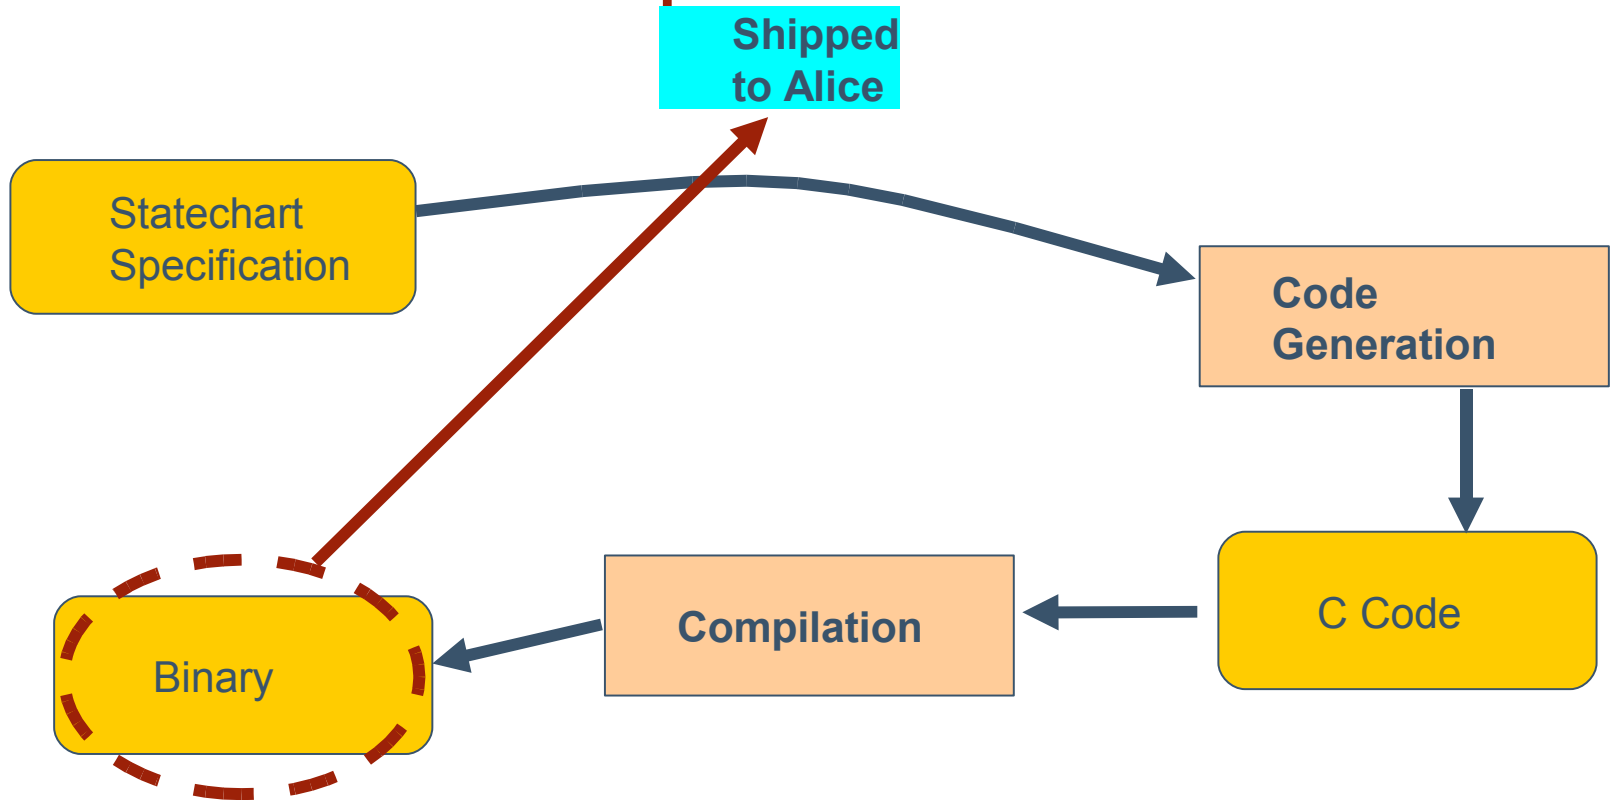
- We already know how to do this for high-level programs
- But binaries are obtained by compiling high-level programs
- Why don't we “**compile**” the **invariants** as well !

The key idea is **invariant compilation**

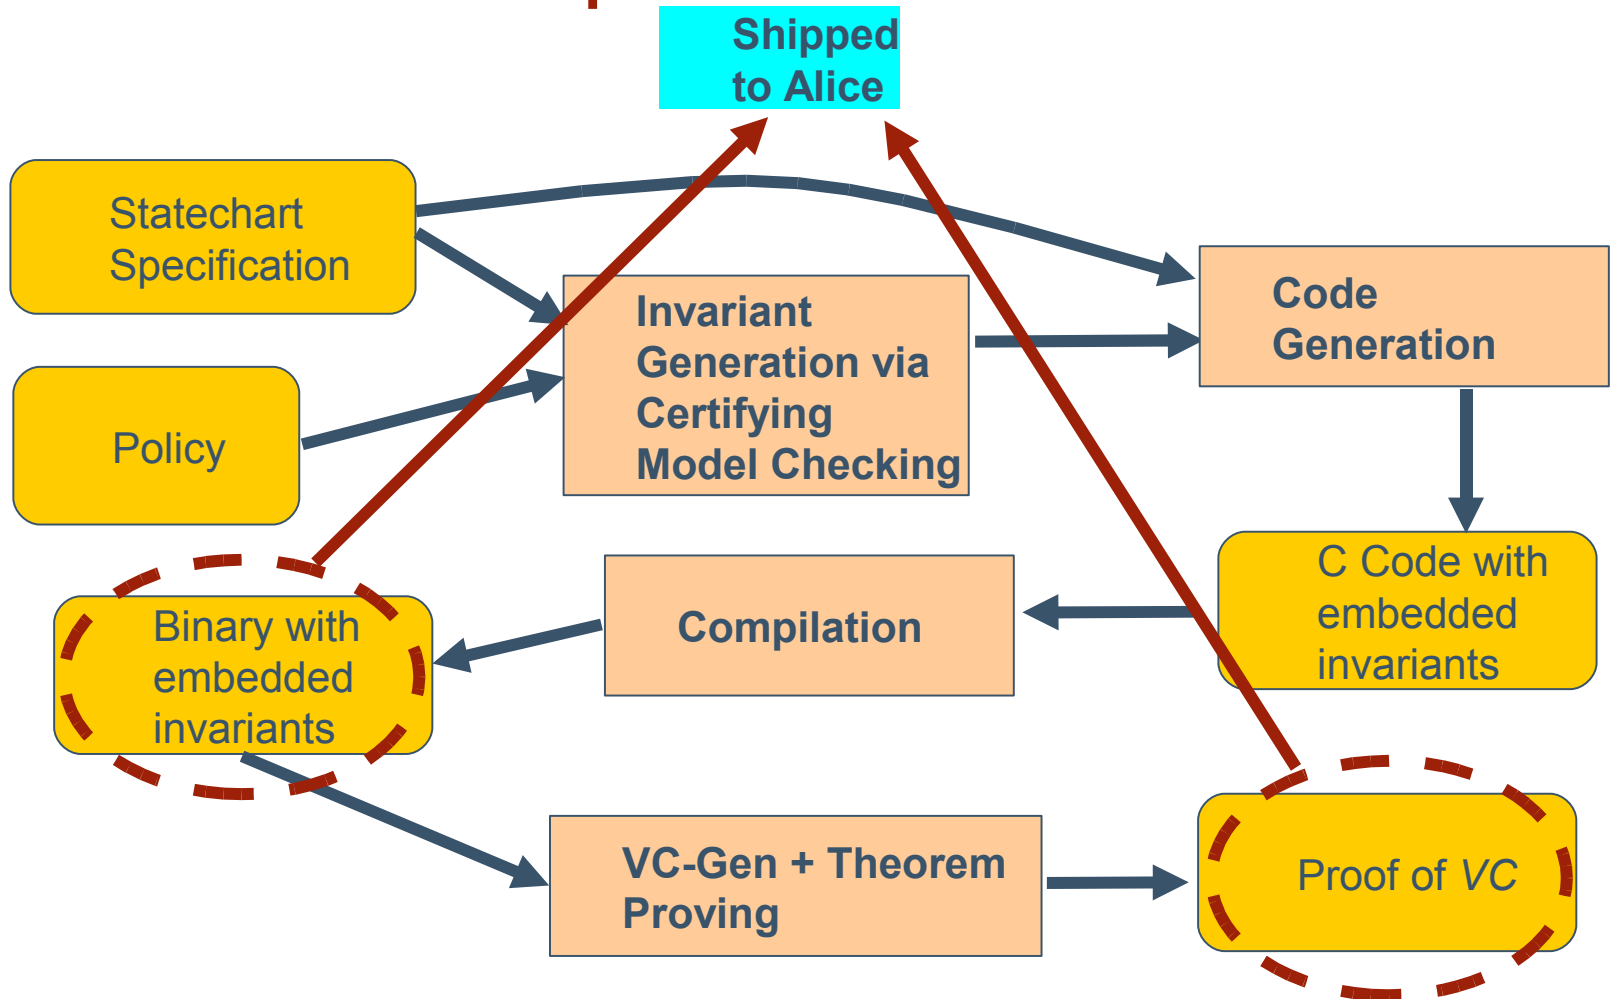
- Independent of the actual programming languages involved
- Need precise semantics for source and target languages, and the translation process



Uncertified Component Binaries



Certified Component Binaries



Experience with an Embedded OS

Micro-C is an lightweight OS for embedded real-time applications

- 6000 lines of C code, supports multi-threading
- uses a “lock” to ensure mutual exclusion

We embedded (manually) appropriate invariants to certify that the OS follows a correct locking convention

Our prototype generated a binary (PowerPC assembly) with embedded invariants as well as the proof of the VC

- used a standard gcc cross compiler
- only default compiler optimizations were allowed

Experience with Embedded OS

Certified the executable of Micro-C Operating System for correct **ordering** of **lock** and **unlock** operations

- Executable size = 335 KB
- Certificate size = 12 KB

Small trusted computing base (**TCB**) source code size

- Translator to SAT = 445 KB
- SAT proof checker = 25 KB

Certifying **model checker** and **compiler** **eliminated** from TCB

- About 20,000 KB source code size
- Over **40 times improvement**

Experience with “TAR”

Program for archiving files

- We used the implementation in the Plan9 OS

Certified that a certain buffer can never be overflowed when the program is executed

- Annotations were added manually
- Applications to certifying software security

Still an application of translating a certified source to a certified binary

Demo : Generating Certified Components

A component that manipulates an integer variable “i”

Policy: The value of “i” should never become negative

The demo consists of three stages:

7. We will generate a certified binary
9. We will modify the binary such that the policy is still obeyed. We will show that the altered binary can still be certified.
11. We will modify the binary such that the policy is violated. We will show that the modified binary can no longer be certified.

Key Results/Conclusion

Component certification involves two broad activities

- invariant generation and invariant compilation

We have a prototypical implementation that was successful in delivering a proof of concept

- Compiled invariants that certify proper locking discipline for an embedded OS and absence of buffer overflows in a commonly used program
- Generated certified binary for a simple illustrative component specification
- Our implementation also supports ranking functions for certifying liveness policies.

Invariants and ranking functions lead to a basis of objective confidence in the behavior of untrusted software

Future Directions

The theoretical foundations have already been laid. There is an urgent need for broader transition to practitioners.

- Organizations who benefit most from this technology
- Medical device industry, auto industry
- Third-party supplied components

“Proof engineering” decisions have to be agreed upon

- This calls for concerted [standardization](#) effort
- Stakeholders who can formulate and enforce formal certification standards

Questions?

<http://www.sei.cmu.edu/staff/chaki>
chaki@sei.cmu.edu