

From Program Verification to Certified Binaries

Angelos Manousaridis Michalis A. Papakyriakou
Nikolaos S. Papaspyrou



National Technical University of Athens
School of Electrical and Computer Engineering
Software Engineering Laboratory

[amanous,mpapakyr,nickie}@softlab.ntua.gr](mailto:{amanous,mpapakyr,nickie}@softlab.ntua.gr)

Logic and Theory of Algorithms
4th Conference on Computability in Europe
Athens, Greece, June 18, 2008

What is this about?

What is this about?

OVERALL RATING: -4 (unacceptable for presentation)

REVIEWER'S CONFIDENCE: 3 (high)

----- REVIEW -----

This short paper **replays the decade old vision of proof-carrying code**, but aiming to increase the level of ambition from simple memory and control-flow properties to **arbitrary program properties**.

(snip)

I was unable to spot any research contributions or novelty in the paper.

(snip)

In summary, this work is much **too preliminary** and is in the current state unacceptable for presentation.

So, what is this all about?

- ▶ A **position paper**, not much of a **research paper**
- ▶ **Goal?** the construction of **certified software**
i.e. that **provably** satisfies its specifications
- ▶ **Why?** the **Holy Grail** of software engineering!
- ▶ **How?** by combining **formal verification**
and **proof-carrying code**

Outline

Introduction

Program verification

Proof-carrying code

Motivation

A Hybrid System

A Motivating Example

Proof-preserving Compilation

Conclusion

Introduction (i)

- ▶ Program verification
 - ▶ aims at formally proving program **correctness**
 - ▶ given a formal **specification** or **property**
 - ▶ **long tradition** (4 decades)
 - ▶ several formal **logics** (e.g. Hoare Logic)
 - ▶ at the **source code** level

Introduction (ii)

- ▶ **Proof-carrying code** (PCC)
 - ▶ **certified binary**: a **value** together with a **proof** that the value satisfies a given specification
 - ▶ relatively **recent** approach (~10 years)
 - ▶ essential in modern **distributed** computer systems
 - ▶ executable code is **transferred** among devices that do not necessarily **trust** one another
 - ▶ at a **lower** level (e.g. machine language)
 - ▶ mainly interested in relatively simple properties: **memory safety** and **control flow**

Introduction (iii)

- ▶ **Type-theoretic** approaches to PCC
 - e.g. Shao *et al.*, POPL 2002, TOPLAS 2005;
Crary and Vanderwaart, ICFP 2002
 - ▶ **arbitrary** program properties
 - ▶ embedding of logic “**formulae as types**”
 - ▶ **proof-preserving** compilation
 - ▶ makes the **proof** a part of the **code**
 - ▶ type (proof) **inference** is undecidable





Motivation

	Program verification	PCC
programmer "friendly"		
high-level proofs		
end-user safety		







Motivation

	Program verification	PCC
programmer "friendly"		
high-level proofs		
end-user safety		







Motivation

	Program verification	PCC
programmer "friendly"		
high-level proofs		
end-user safety		

Motivation

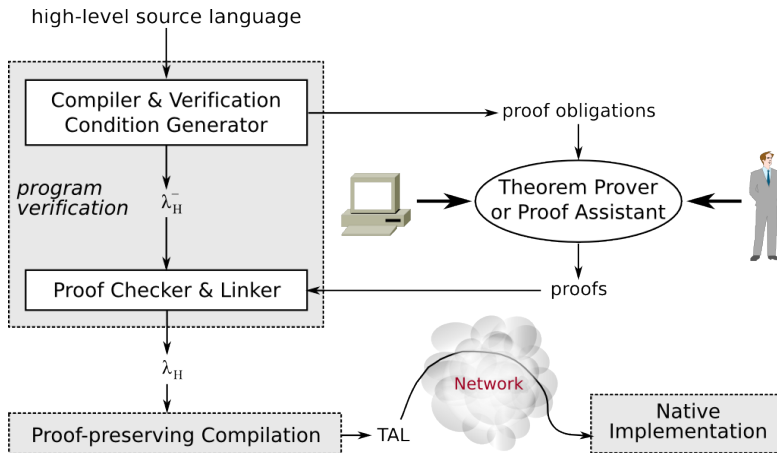
	Program verification	PCC
programmer "friendly"		
high-level proofs		
end-user safety		

Motivation

	Program verification	PCC
programmer "friendly"		
high-level proofs		
end-user safety		

Can we write programs in a **high-level** language,
provide **correctness proofs** for them
and then **compile** them
to **provably** correct **executable** code?

A hybrid system



Integer square root (i)

- ▶ Given $n \in \mathbb{N}$, find the **greatest** $r \in \mathbb{N}$ such that $r^2 \leq n$

Integer square root (i)

- ▶ Given $n \in \mathbb{N}$, find the **greatest** $r \in \mathbb{N}$ such that $r^2 \leq n$

```
int root (int n) {  
    int y = 0;  
  
    while ((y+1)*(y+1) <= n) y++;  
    return y;  
}
```


Integer square root (i)

- ▶ Given $n \in \mathbb{N}$, find the **greatest** $r \in \mathbb{N}$ such that $r^2 \leq n$

//@ predicate leRoot(int r, int x) { $r \geq 0 \wedge r^2 \leq x$ }

//@ predicate isRoot(int r, int x) { leRoot(r, x) $\wedge (r + 1)^2 > x$ }

```
int root (int n) {  
    int y = 0;  
  
    while ((y+1)*(y+1) <= n) y++;  
    return y;  
}
```

Integer square root (i)

- ▶ Given $n \in \mathbb{N}$, find the **greatest** $r \in \mathbb{N}$ such that $r^2 \leq n$

```
//@ predicate leRoot(int r, int x) { r ≥ 0 ∧ r2 ≤ x }
//@ predicate isRoot(int r, int x) { leRoot(r, x) ∧ (r + 1)2 > x }

/*@ requires n ≥ 0
    @ ensures  isRoot(\result, n)
    @*/
int root (int n) {
    int y = 0;

    while ((y+1)*(y+1) <= n) y++;
    return y;
}
```

Integer square root (i)

- ▶ Given $n \in \mathbb{N}$, find the **greatest** $r \in \mathbb{N}$ such that $r^2 \leq n$

```
//@ predicate leRoot(int r, int x) { r ≥ 0 ∧ r2 ≤ x }
//@ predicate isRoot(int r, int x) { leRoot(r, x) ∧ (r + 1)2 > x }

/*@ requires n ≥ 0
    @ ensures  isRoot(\result, n)
    @*/
int root (int n) {
    int y = 0;
    //@ invariant leRoot(y, n)
    while ((y+1)*(y+1) <= n) y++;
    return y;
}
```

Integer square root (ii)

Proof obligations

in Why/Caduceus

1. $\forall n \in \mathbb{Z}.$

$$n \geq 0 \Rightarrow \text{leRoot}(0, n)$$

2. $\forall n, y \in \mathbb{Z}.$

$$n \geq 0 \wedge \text{leRoot}(y, n) \wedge (y + 1)^2 \leq n \Rightarrow \\ \text{leRoot}(y + 1, n)$$

3. $\forall n, y \in \mathbb{Z}.$

$$n \geq 0 \wedge \text{leRoot}(y, n) \wedge (y + 1)^2 > n \Rightarrow \\ \text{isRoot}(y, n)$$

They are all automatically discharged, using the definitions of **leRoot** and **isRoot**

Integer square root (ii)

Proof obligations

translation to \overline{H}

```
root  ▷  ∀n:Z. ∀n*:(n ≥ 0). sint n → ∃x:Z. ∃x*:isRoot x n. sint x
      =  poly n:Z. poly n*:(n ≥ 0). lambda n:sint n.
        (fix loop:∀y:Z. ∀y*:leRoot y n. sint y →
          ∃x:Z. ∃x*:isRoot x n. sint x.
          poly y:Z. poly y*:leRoot y n. lambda y:sint y.
            if [♣, ♣] ((y + cint [1])2 > n,
              p1* . pack (y, pack (♣, y) as ∃y*:isRoot y n. sint y) as
                ∃x:Z. ∃x*:isRoot x n. sint x,
              p2* . loop [y + 1] [♣] (y + cint [1])))
        [0] [♣] cint [0]
```

Integer square root (iii)

Proof obligations

from \bar{H} to H

```
root  ▷  ∀n:Z. ∀n*:(n ≥ 0). sint n → ∃x:Z. ∃x*:isRoot x n. sint x
=      poly n:Z. poly n*:(n ≥ 0). lambda n:sint n.
      (fix loop:∀y:Z. ∀y*:leRoot y n. sint y →
        ∃x:Z. ∃x*:isRoot x n. sint x.
        poly y:Z. poly y*:leRoot y n. lambda y:sint y.
        if [decidable ((y + 1)2 > n), gtDecidablePrf (y + 1)2 n] (
          (y + cint [1])2 > n,
          p1* . pack (y, pack (conj y* p1*, y) as
                     ∃y*:isRoot y n. sint y) as
                     ∃x:Z. ∃x*:isRoot x n. sint x,
          p2* . loop [y + 1] [conj (Zplus_ge_compat y 0 1 0
                                   (proj1 y*) (geDecidablePrf 1 0))
                                   (Znot_gt_le (y + 1)2 n p2*)]
                     (y + cint [1])))
        [0] [conj (Z_ge_refl 0) (Zge_le n 0 n*)] cint [0])
```

Proof-preserving compilation (i)

Continuation passing style (CPS) from H to K

- ▶ Functions **do not return**
- ▶ One more parameter: the **continuation**
- ▶ **Jumps** instead of **calls**
- ▶ Control flow is **explicit**
- ▶ **Optimizing** transformations can be applied
- ▶ **~20 lines** for the square root example

Proof-preserving compilation (ii)

Closure conversion

from K to C

- ▶ Functions only use **local data**
- ▶ One more parameter: the **closure**
- ▶ More **optimizing** transformations can be applied
- ▶ **~200 lines** for the square root example

Proof-preserving compilation (iii)

Hoisting

from C to A

- ▶ All functions become **top-level** blocks
- ▶ Memory allocation is **explicit**
- ▶ **~200 lines** for the square root example

Proof-preserving compilation (iv)

Typed assembly language (TAL) from A to TAL

- ▶ RISC
- ▶ We assume **infinite registers**
 - ▶ no **spilling** phase
 - ▶ trivial **register allocation**
- ▶ We assume a **garbage collector**
- ▶ **~500 lines** for the square root example

Proof-preserving compilation (v)

Beyond compilation

- ▶ **Low Level Virtual Machine** (LLVM)
- ▶ Direct translation from TAL to LLVM
- ▶ Direct translations from LLVM to native code for many architectures

x86, x86-64, PowerPC 32/64, ARM, Thumb, IA-64, Alpha, SPARC, MIPS, CellSPU

Conclusion

- ▶ Long, long way to go...

Conclusion

- ▶ Long, long way to go...
- ▶ Most problems on the source level
 - ▶ what language(s)?
 - ▶ what logic(s)?
 - ▶ can programmers easily prove things?
 - ▶ scalability?

Conclusion

- ▶ Long, long way to go...
- ▶ Most problems on the source level
 - ▶ what language(s)?
 - ▶ what logic(s)?
 - ▶ can programmers easily prove things?
 - ▶ scalability?
- ▶ Efficient representations for certified binaries?

Conclusion

- ▶ Long, long way to go...
- ▶ Most problems on the source level
 - ▶ what language(s)?
 - ▶ what logic(s)?
 - ▶ can programmers easily prove things?
 - ▶ scalability?
- ▶ Efficient representations for certified binaries?
- ▶ Can modern compilers be proof-preserving?

Conclusion

Thank you!

- ▶ Long, long way to go...
- ▶ Most problems on the source level
 - ▶ what language(s)?
 - ▶ what logic(s)?
 - ▶ can programmers easily prove things?
 - ▶ scalability?
- ▶ Efficient representations for certified binaries?
- ▶ Can modern compilers be proof-preserving?