

PROOF-CARRYING-CODE

Hans-Wolfgang Loidl

LFE Theoretische Informatik, Institut für Informatik,
Ludwig-Maximilians Universität, München

June 19, 2007

- ① MOTIVATION
- ② BASIC CONCEPTS
- ③ AN EXAMPLE: CCURED
- ④ MAIN CHALLENGES
 - Certificate Size
 - Performance
 - Size of the TCB
- ⑤ SUMMARY

MOTIVATION

Downloading software over the network is nowadays common-place.

But who says that the software does what it promises to do?

Who protects the consumer from malicious software or other undesirable side-effects?

MOTIVATION

Downloading software over the network is nowadays common-place.

But who says that the software does what it promises to do?

Who protects the consumer from malicious software or other undesirable side-effects?

⇒ **Mechanisms for ensuring that a program is “well-behaved” are needed.**

AUTHENTICATION FOR MOBILE CODE

The main mechanisms used nowadays are based on **authentication**.

Java: define safety policies to control the level of safety; managed through cryptographic signatures on the code.

Windows: Microsoft's Authenticode attaches cryptographic signatures to the code; more or less compulsory in Windows XP for drivers.

AUTHENTICATION FOR MOBILE CODE

The main mechanisms used nowadays are based on **authentication**.

Java: define safety policies to control the level of safety; managed through cryptographic signatures on the code.

Windows: Microsoft's Authenticode attaches cryptographic signatures to the code; more or less compulsory in Windows XP for drivers.

But, all these mechanisms say nothing about the code, only about the supplier of the code!

WHOM DO YOU TRUST COMPLETELY?



MAYBE THAT'S NOT SUCH A GOOD IDEA!

Microsoft Security Bulletin MS01-017

Who should read this bulletin: All customers using Microsoft® products.

Technical description: In mid-March 2001, VeriSign, Inc., advised Microsoft that on January 29 and 30, 2001, it issued two VeriSign Class 3 code-signing digital certificates to an individual who fraudulently claimed to be a Microsoft employee. ...

Impact of vulnerability: Attacker could digitally sign code using the name "Microsoft Corporation".

PROOF-CARRYING-CODE (PCC): THE IDEA

Goal: Safe execution of untrusted code.

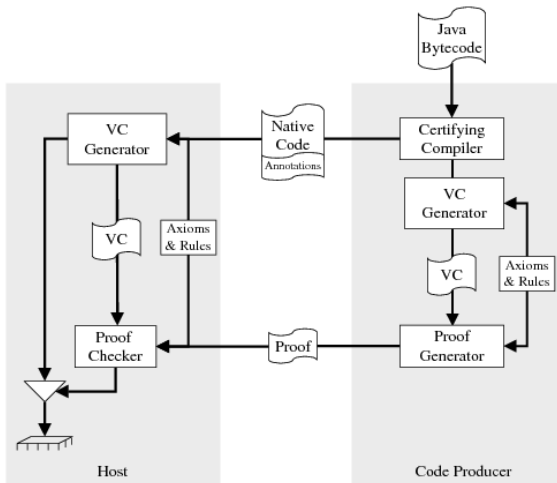
PCC is a software mechanism that allows a host system to determine with certainty that it is safe to execute a program supplied by an untrusted source.

Method: Together with the code, a *certificate* describing its behaviour is sent.

This certificate is a condensed form of a formal proof of this behaviour.

Before execution, the consumer can check the behaviour, by running the proof against the program.

A PCC ARCHITECTURE



PROGRAM VERIFICATION TECHNIQUES

Many techniques for PCC come from the area of **program verification**. Main differences:

General program verification

- is trying to **verify good behaviour** (correctness).
- is usually interactive
- requires at least programmer annotations as invariants to the program

PROGRAM VERIFICATION TECHNIQUES

Many techniques for PCC come from the area of **program verification**. Main differences:

General program verification

- is trying to **verify good behaviour** (correctness).
- is usually interactive
- requires at least programmer annotations as invariants to the program

PCC

- is trying to **falsify bad behaviour**
- must be automatic
- may be based on inferred information from the high-level

PROGRAM VERIFICATION TECHNIQUES

Many techniques for PCC come from the area of **program verification**. Main differences:

General program verification

- is trying to **verify good behaviour** (correctness).
- is usually interactive
- requires at least programmer annotations as invariants to the program

PCC

- is trying to **falsify bad behaviour**
- must be automatic
- may be based on inferred information from the high-level

Observation: Checking a proof is much simpler than creating one

AN EXAMPLE: CCURED

CCured is a system for checking **pointer-safety** of C programs, developed by the group of George Necula at Berkeley.

Uses a hybrid mechanism of static type checking and run-time checks.

Goal: Prove pointer safety statically, where possible, and minimise required run-time checks.

A ROADMAP TO A PCC INFRASTRUCTURE

Task of the infrastructure: **Certify** that the **execution** of the program is **well-behaved**.

Several steps to build the infrastructure:

- Formalise **execution** as an **operational semantics** of the language.
- Formalise **well-behaved** as a **safety policy** (type-system)
- **Certify** safety by producing a proof-term (or similar).

THE CORE LANGUAGE

Mini-C language:

$$\begin{aligned} e &::= x \mid n \mid e_1 \text{ op } e_2 \mid (\tau)e \mid e_1 \oplus e_2 \mid !e \\ c &::= \text{skip} \mid c_1; c_2 \mid e_1 := e_2 \end{aligned}$$

Types: standard C types with extension for **pointers into arrays** and **dynamic types**.

Efficient type inference is possible and demonstrated for this type system.

THE CCURED TYPE SYSTEM: POINTERS

C contains 2 evil pointer operations: arithmetic and casts.

The type system distinguishes between 3 kinds of pointers:

- **Safe pointers**: no arithmetic or casts; represented as an address
- **Sequence pointers**: arithmetic but **no casts**; represented as a region
- **Dynamic pointers**: **casts**, all bets are off! represented as a region

EXAMPLE PROGRAM

Sum over an array of boxed integers:

```
int acc; /* accumulator */ int **p; // elem ptr
int **a; /* array */      int i;  // index
int *e;  /* unboxer */
acc = 0;
for (i=0; i<100; i++) {
    p = a + i;           // ptr arithm
    e = *p;             // read elem
    while ((int)e % 2 == 0) { // check tag
        e = *(int **)e; // unbox
    }
    acc += ((int)e >> 1); // strip tag
}
```

EXAMPLE PROGRAM

Sum over an array of boxed integers:

```
int acc; /* accumulator */ int **p; // elem ptr
int **a; /* array */          int i;  // index
int *e;  /* unboxer */
acc = 0;
for (i=0; i<100; i++) {
    p = a + i;                // ptr arithm
    e = *p;                   // read elem
    while ((int)e % 2 == 0) { // check tag
        e = *(int **)e;      // unbox
    }
    acc += ((int)e >> 1);    // strip tag
}
```

a and p point into an array with elems of type int *

EXAMPLE PROGRAM

Sum over an array of boxed integers:

```
int acc; /* accumulator */ int **p; // elem ptr
int **a; /* array */          int i; // index
int *e; /* unboxer */
acc = 0;
for (i=0; i<100; i++) {
    p = a + i; // ptr arithm
    e = *p; // read elem
    while ((int)e % 2 == 0) { // check tag
        e = *(int **)e; // unbox
    }
    acc += ((int)e >> 1); // strip tag
}
```

a is subject to pointer arithm (“**sequence pointer**”)

⇒ check for out of bounds

EXAMPLE PROGRAM

Sum over an array of boxed integers:

```
int acc; /* accumulator */ int **p; // elem ptr
int **a; /* array */      int i;  // index
int *e;  /* unboxer */
acc = 0;
for (i=0; i<100; i++) {
    p = a + i;                // ptr arithm
    e = *p;                  // read elem
    while ((int)e % 2 == 0) { // check tag
        e = *(int **)e;     // unbox
    }
    acc += ((int)e >> 1);    // strip tag
}
```

p has no arithmetic (“safe pointer”)

⇒ no bounds check needed

EXAMPLE PROGRAM

Sum over an array of boxed integers:

```
int acc; /* accumulator */ int **p; // elem ptr
int **a; /* array */      int i;  // index
int *e;  /* unboxer */
acc = 0;
for (i=0; i<100; i++) {
    p = a + i;           // ptr arithm
    e = *p;             // read elem
    while ((int)e % 2 == 0) { // check tag
        e = *(int **)e; // unbox
    }
    acc += ((int)e >> 1); // strip tag
}
```

e is subject to a type cast (“dynamic pointer”)

⇒ nothing known about underlying type

OPERATIONAL SEMANTICS

The value of an integer, or a safe pointer is an integer n ; the value of a sequence or dynamic pointer is a **home**, modelled as a pair $\mathbb{N} \times \mathbb{N}$ of start address and offset.

$$v ::= n \mid \langle h, n \rangle$$

OPERATIONAL SEMANTICS

The value of an integer, or a safe pointer is an integer n ; the value of a sequence or dynamic pointer is a **home**, modelled as a pair $\mathbb{N} \times \mathbb{N}$ of start address and offset.

$$v ::= n \mid \langle h, n \rangle$$

Each home is tagged as being an integer or a pointer, and has an associated **kind** and **size** functions. The semantic domain for pointers:

$$\begin{aligned} \llbracket \text{int} \rrbracket_H &= \mathbb{N} \\ \llbracket \text{DYNAMIC} \rrbracket_H &= \{ \langle h, n \rangle \mid h \in H \wedge (h = 0 \vee \text{kind}(h) = \text{untyped}) \} \\ \llbracket \tau \text{ ref SEQ} \rrbracket_H &= \{ \langle h, n \rangle \mid h \in H \wedge (h = 0 \vee \text{kind}(h) = \text{typed}(\tau)) \} \\ \llbracket \tau \text{ ref SAFE} \rrbracket_H &= \{ h + i \mid h \in H \wedge 0 \leq i \leq \text{size}(h) \wedge \\ &\quad (h = 0 \vee \text{kind}(h) = \text{typed}(\tau)) \} \end{aligned}$$

OPERATIONAL SEMANTICS (POINTERS)

$$\frac{\Sigma, M \vdash e_1 \Downarrow \langle h, n_1 \rangle \quad \Sigma, M \vdash e_2 \Downarrow n_2}{\Sigma, M \vdash e_1 \oplus e_2 \Downarrow \langle h_1, n_1 + n_2 \rangle} \quad (\text{POINTER ARITHM})$$

$$\frac{\Sigma, M \vdash e \Downarrow \langle h, n \rangle}{\Sigma, M \vdash (\text{int})e \Downarrow h + n} \quad (\text{CASTTOINT})$$

$$\frac{\Sigma, M \vdash e \Downarrow n}{\Sigma, M \vdash (\tau \text{ ref SEQ})e \Downarrow \langle 0, n \rangle} \quad (\text{CASTTOSEQ})$$

$$\frac{\Sigma, M \vdash e \Downarrow \langle h, n \rangle \quad \mathbf{0 \leq n \leq \text{size}(h)}}{\Sigma, M \vdash (\tau \text{ ref SAFE})e \Downarrow h + n} \quad (\text{CASTTOSAFE})$$

OPERATIONAL SEMANTICS (READ OPERATIONS)

Two kinds of reads, with different obligations for run-time checks:

$$\frac{\Sigma, M \vdash e \Downarrow n \quad \mathbf{n \neq 0}}{\Sigma, M \vdash !e \Downarrow M(n)} \quad (\text{SAFERD})$$

$$\frac{\Sigma, M \vdash e \Downarrow \langle h, n \rangle \quad \mathbf{h \neq 0} \quad \mathbf{0 \leq n \leq \text{size}(h)}}{\Sigma, M \vdash !e \Downarrow M(h + n)} \quad (\text{DYNRD})$$

$$\frac{\Sigma, M \vdash e_1 \Downarrow n \quad \mathbf{n \neq 0} \quad \Sigma, M \vdash e_2 \Downarrow v}{\Sigma, M \vdash e_1 := e_2 \Downarrow M(n \mapsto v)} \quad (\text{SAFEWR})$$

$$\frac{\Sigma, M \vdash e_1 \Downarrow \langle h, n \rangle \quad \mathbf{h \neq 0} \quad \mathbf{0 \leq n \leq \text{size}(h)} \quad \Sigma, M \vdash e_2 \Downarrow v}{\Sigma, M \vdash e_1 := e_2 \Downarrow M(h + n \mapsto v)} \quad (\text{DYNWR})$$

THE CCURED TYPE SYSTEM: RULES

The type system keeps track of the kind of pointers.
Rules for converting pointers:

$$\overline{\tau \leq \tau}$$

$$\overline{\tau \leq \text{int}}$$

$$\overline{\text{int} \leq \tau \text{ ref SEQ}}$$

$$\overline{\text{int} \leq \text{DYNAMIC}}$$

$$\overline{\tau \text{ ref SEQ} \leq \tau \text{ ref SAFE}}$$

TYPING RULES FOR COMMANDS

$\Gamma \vdash c$ means, command c is well-typed.

$\Gamma \vdash e : \tau$ means, expression e has type τ .

$$\frac{}{\Gamma \vdash \text{skip}}$$

$$\frac{\Gamma \vdash c_1 \quad \Gamma \vdash c_2}{\Gamma \vdash c_1; c_2}$$

$$\frac{\Gamma \vdash e : \tau \text{ ref SAFE} \quad \Gamma \vdash e' : \tau}{\Gamma \vdash e := e'}$$

$$\frac{\Gamma \vdash e : \text{DYNAMIC} \quad \Gamma \vdash e' : \text{DYNAMIC}}{\Gamma \vdash e := e'}$$

TYPING RULES FOR EXPRESSIONS

$$\frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau}$$

$$\frac{\Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 \text{ op } e_2 : \text{int}}$$

$$\frac{\Gamma \vdash e : \tau' \quad \tau' \leq \tau}{\Gamma \vdash (\tau)e : \tau}$$

$$\frac{}{\Gamma \vdash (\tau \text{ ref SAFE})0 : \tau \text{ ref SAFE}}$$

$$\frac{\Gamma \vdash e_1 : \tau \text{ ref SEQ} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 \oplus e_2 : \tau \text{ ref SEQ}}$$

$$\frac{\Gamma \vdash e_1 : \text{DYNAMIC} \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1 \oplus e_2 : \text{DYNAMIC}}$$

$$\frac{\Gamma \vdash e : \tau \text{ ref SAFE}}{\Gamma \vdash !e : \tau}$$

$$\frac{\Gamma \vdash e : \text{DYNAMIC}}{\Gamma \vdash !e : \text{DYNAMIC}}$$

SAFETY POLICY

The **safety policy** states, that at all times in the execution, the contents of each memory address must correspond to the typing constraints of the home to which it belongs.

Formally, the following predicate must be fulfilled at all times

$$\begin{aligned}
 WF(M_H) \equiv & \forall h \in H^*. \forall i \in \mathbb{N}. 0 \leq i < \text{size}(h) \Rightarrow \\
 & (\text{kind}(h) = \text{untyped} \Rightarrow M(h+i) \in \parallel \text{DYNAMIC} \parallel_H \wedge \\
 & \text{kind}(h) = \text{typed}(\tau) \Rightarrow M(h+i) \in \parallel \tau \parallel_H
 \end{aligned}$$

We can prove that this property is preserved by all rules in the type system.

THEOREMS

We separate run-time failure from rightful termination like this:

$\Sigma, M_H \vdash e \Downarrow \textit{CheckFailed}$ means a run-time check failed during the execution of expression e .

THEOREMS

We separate run-time failure from rightful termination like this:
 $\Sigma, M_H \vdash e \Downarrow \text{CheckFailed}$ means a run-time check failed during the execution of expression e .

THEOREM

(Progress and type preservation) If $\Gamma \vdash e : \tau$ and $\Sigma \in \llbracket \Gamma \rrbracket_H$ and $WF(M_H)$, then either $\Sigma, M_H \vdash e \Downarrow \text{CheckFailed}$ or $\Sigma, M_H \vdash e \Downarrow v$ and $v \in \llbracket \tau \rrbracket_H$.

THEOREMS

$\Sigma, M_H \vdash c \implies \textit{CheckFailed}$ means a run-time check failed during the execution of command c .

THEOREMS

$\Sigma, M_H \vdash c \implies \text{CheckFailed}$ means a run-time check failed during the execution of command c .

THEOREM

(Progress for commands) If $\Gamma \vdash c$ and $\Sigma \in \|\Gamma\|_h$ and $WF(M_H)$ then either $\Sigma, M_H \vdash c \implies \text{CheckFailed}$ or $\Sigma, M_H \vdash c \implies M'_H$ and M'_H is well-formed.

MAIN RESULTS

- An efficient inference algorithm attaches `ref SEQ`, `ref SAFE`, `DYNAMIC` annotations to plain C code.
- Most of the checks can be done statically.
- The performance overhead of the remaining run-time checks is moderate: 0–150%
- Purely dynamic checks would incur a performance overhead of factors 6–20
- Several array bounds bugs discovered in SPECINT95




FURTHER READING



CCured: Type-Safe Retrofitting of Legacy Code, in POPL'02
— ACM Symposium on Principles of Programming Languages,
2002. Online Demo at

<http://manju.cs.berkeley.edu/ccured/web/index.html>.

FURTHER READING

-  George Necula, *Proof-carrying code* in POPL'97 —
Symposium on Principles of Programming Languages, Paris,
France, 1997.
http://raw.cs.berkeley.edu/Papers/pcc_popl97.ps
-  George Necula, *Proof-Carrying Code: Design and
Implementation* in Proof and System Reliability,
Springer-Verlag, 2002.
<http://raw.cs.berkeley.edu/Papers/marktoberdorf.pdf>
-  *CCured Demo*,
<http://manju.cs.berkeley.edu/ccured/web/index.html>

MAIN CHALLENGES OF PCC

PCC is a very powerful mechanism. Coming up with an efficient implementation of such a mechanism is a challenging task.

The main problems are

- Certificate size
- Performance of validation
- Size of the trusted code base (TCB)

CERTIFICATE SIZE

A certificate is a formal proof, and can be encoded as e.g. LF Term.

BUT: such proof terms include a lot of repetition
⇒ huge certificates

Approaches to reduce certificate size:

- Compress the general proof term and do reconstruction on the consumer side
- Transmit only hints in the certificate (oracle strings)
- Embed the proving infrastructure into a theorem prover and use its tactic language

PERFORMANCE

Even though validation is fast compared to proof generation, it is on the critical path of using remote code
⇒ performance of the validation is crucial for the acceptance of PCC.

Approaches:

- Write your own specialised proof-checker (for a specific domain)
- Use hooks of a general proof-checker, but replace components with more efficient routines, e.g. arithmetic

SIZE OF THE TRUSTED CODE BASE (TCB)

The PCC architecture relies on the correctness of components such as VC-generation and validation.

But these components are complex and implementation is error-prone.

Approaches for reducing size of TCB:

- Use proven/established software
- Build everything up from basics **foundational PCC** (Appel)

FOUNDATIONAL PCC

Philosophy of Foundational PCC: Minimise the “trusted code base”, i.e. the software that needs to be trusted.

Approach of Foundational PCC: Define safety policy directly on the **operational semantics** of the code.

Certificates are proofs over the operational semantics.

FOUNDATIONAL PCC

Philosophy of Foundational PCC: Minimise the “trusted code base”, i.e. the software that needs to be trusted.

Approach of Foundational PCC: Define safety policy directly on the **operational semantics** of the code.

Certificates are proofs over the operational semantics.

Pros and cons:

- 😊 **more flexible:** not restricted to a particular type system as the language in which the proofs are phrased;
- 😊 **more secure:** no reliance on VCG.
- 😞 **larger proofs**

CONVENTIONAL VS FOUNDATIONAL PCC

Re-examine the logic for memory safety, eg.

$$\frac{m \vdash e : \tau \text{ list} \quad e \neq 0}{m \vdash e : \text{addr} \wedge m \vdash e + 4 : \text{addr} \wedge m \vdash \text{sel}(m, e) : \tau \wedge m \vdash \text{sel}(m, e + 4) : \tau \text{ list}} \quad (\text{LISTELIM})$$

The rule has **built-in knowledge about the type-system**, in this case representing the data layout of the compiler (“*Type specialised PCC*”) \implies dangerous if soundness of the logic is not checked mechanically!

LOGIC RULES IN FOUNDATIONAL PCC

In foundational PCC the rules work on the operational semantics:

$$\frac{m \models e : \tau \text{ list} \quad e \neq 0}{m \models e : \text{addr} \wedge m \models e + 4 : \text{addr} \wedge m \models \text{sel}(m, e) : \tau \wedge m \models \text{sel}(m, e + 4) : \tau \text{ list}} \quad (\text{LISTELIM})$$

This looks similar to the previous rule but has a very different meaning: \models is a predicate over the formal model of the computation, and the above rule can be proven as a lemma, \vdash is an encoding of a type-system on top of the operational semantics and thus needs a **soundness proof**.

EXAMPLE: SPECIFYING SAFE MEMORY ACCESS

To specify safety, the operational semantics is written in such a way, that it gets stuck whenever the safety condition is violated.

EXAMPLE: SPECIFYING SAFE MEMORY ACCESS

To specify safety, the operational semantics is written in such a way, that it gets stuck whenever the safety condition is violated.

Example: operational semantics on assembler code.

Safety policy: “only readable addresses are loaded”.

Define a predicate: $readable(x) \equiv 0 \leq x \leq 1000$

EXAMPLE: SPECIFYING SAFE MEMORY ACCESS

To specify safety, the operational semantics is written in such a way, that it gets stuck whenever the safety condition is violated.

Example: operational semantics on assembler code.

Safety policy: “only readable addresses are loaded”.

Define a predicate: $readable(x) \equiv 0 \leq x \leq 1000$

The semantics of a load operation $LD\ r_i, c(r_j)$ is now written as follows:

$$\begin{aligned}
 load(i, j, c) &\equiv \lambda r\ m\ r'\ m'. \\
 &\quad r'(i) = m(r(j) + c) \wedge readable(r(j) + c) \wedge \\
 &\quad (\forall x \neq i. r'(x) = r(x)) \wedge m' = m
 \end{aligned}$$

EXAMPLE: SPECIFYING SAFE MEMORY ACCESS

To specify safety, the operational semantics is written in such a way, that it gets stuck whenever the safety condition is violated.

Example: operational semantics on assembler code.

Safety policy: “only readable addresses are loaded”.

Define a predicate: $readable(x) \equiv 0 \leq x \leq 1000$

The semantics of a load operation $LD\ r_i, c(r_j)$ is now written as follows:

$$load(i, j, c) \equiv \lambda r\ m\ r'\ m'. \\ r'(i) = m(r(j) + c) \wedge readable(r(j) + c) \wedge \\ (\forall x \neq i. r'(x) = r(x)) \wedge m' = m$$

Note: the clause for nothing else changes, quickly becomes awkward when doing these proofs

⇒ Separation Logic (Reynolds'02) tackles this problem.

FURTHER READING



Andrew Appel, *Foundational Proof-Carrying Code* in LICS'01
— Symposium on Logic in Computer Science, 2001.

<http://www.cs.princeton.edu/~appel/papers/fpcc.pdf>

SUMMARY

PCC is a powerful, general mechanism for providing safety guarantees for mobile code.

It provides these guarantees without resorting to a trust relationship.

It uses techniques from the areas of type-systems, program verification and logics.

It is a very active research area at the moment.

PCC reading list:

<http://www.tcs.ifi.lmu.de/~hwloidl/PCC/reading.html>