



COLLÈGE
DE FRANCE
—1530—

Language-based software security, fifth cours

Typing and security

Xavier Leroy

2022-04-07

Collège de France, chair of software sciences

`xavier.leroy@college-de-france.fr`

Earlier, we saw that run-time safety is necessary for software security.

In this lecture, we study **strong typing** of programming languages

- as the primary mean to guarantee run-time safety;
- as a mean to obtain security guarantees that go beyond safety.

Typing in programming languages

Folk wisdom:

Don't compare apples and oranges.

On n'additionne pas des choux et des carottes.

Physicist's wisdom: dimensional analysis.

$d = v.t$ $\text{dist} = \text{dist}$ homogeneous, possibly correct


$d = v/t$ $\text{dist} \neq \text{dist} \cdot \text{temps}^{-2}$ not homogeneous, must be wrong

Types to aid compilation

As early as FORTRAN I (1957), type declarations determine the in-memory representation of data and guide machine code generation.

```
float t[10][20];    int i,j;    float x;  
x = x + i * t[i][j];
```

In-memory representation of `t`: 200×4 bytes



Evaluation of `x`:

$$x +^{\text{float}} (\text{floatofint}(i) \times^{\text{float}} \text{load}(\text{float}, t + (i \times 10 + j) \times 4))$$

Types as a modeling device

A method is proposed for the representation in a computer of complex structured objects, and for their manipulation by a program written in a general purpose language, which is here assumed to be an extension of ALGOL 60.

(C.A.R. Hoare, Record handling, 1965)

Introduces **records** with named, typed fields

```
coloredpoint = { x: float; y: float; c: color }
```

and **references** to point from one record to another

```
intlist = { head: int; tail: ↑intlist }
```

Modeling with records and references

Hoare's example: representing a set of persons and their family relationships.

```
record class person
begin
  integer date of birth;
  Boolean male;
  reference father, mother,
             youngest offspring, elder sibling (person)
end;
```

Types to prevent programming errors

(C.A.R. Hoare, *Notes on data structuring*, 1970).

[T]he use of a high-level language [...] significantly reduces the scope for programming error.

In machine code programming it is all too easy to make stupid mistakes, such as using fixed point addition on floating point numbers, performing arithmetic operations on Boolean markers, or allowing modified addresses to go out of range.

Types to prevent programming errors

(C.A.R. Hoare, *Notes on data structuring*, 1970).

When using a high-level language, such errors may be prevented by three means:

- 1. Errors involving the use of the wrong arithmetic instructions are logically impossible; no program expressed, for example in ALGOL, could ever cause such erroneous code to be generated.*
- 2. Errors like performing arithmetic operations on Boolean markers will be immediately detected by a compiler, and can never cause trouble in an executable program.*
- 3. Errors like the use of a subscript out of range can be detected by runtime checks on the ranges of array subscripts.*

Types to prevent programming errors

The view of typing as a program verification that avoids errors appears in the PhD thesis of J. H. Morris (1968, MIT):

*We shall now introduce a **type system** which, in effect, singles out a decidable subset of those [expressions] that are safe; i.e., cannot give rise to ERRORS.*

This will disqualify certain [expressions] which do not, in fact, cause ERRORS and thus reduce the expressive power of the language.

Morris' type system \approx simple types for the λ -calculus.

Typing and run-time safety

In 1978, R. Milner, in his article *A theory of type polymorphism in programming*, states an essential property of a type system:

an expression (or program) with a legal type assignment cannot “go wrong”

Since then, this property has been used as the characterization of a sound type system.

What does *going wrong* means?

Milner (1978) gives his language a denotational semantics based on Scott domains:

$$\mathcal{E} : Expr \rightarrow Env \rightarrow V$$

where

$$V = (Int + \dots + (V \rightarrow V) + \{\text{wrong}\})_{\perp}$$

`wrong` is the denotation of nonsensical expressions such as

1 2 (the integer 1 used as if it were a function)

1 + ($\lambda x.x$) (a function used as if it were an integer)

An operational semantics based on reductions

Terms $a ::= n \mid x \mid \lambda x.a \mid a_1 a_2 \mid \text{add}$

Values $v ::= n \mid \lambda x.v \mid \text{add} \mid \text{add } v$

$$\begin{array}{c} (\lambda x.a) v \rightarrow a[x \leftarrow v] \\ \\ \frac{a \rightarrow a'}{a b \rightarrow a' b} \end{array} \qquad \begin{array}{c} \frac{n = n_1 + n_2}{\text{add } n_1 n_2 \rightarrow n} \\ \\ \frac{b \rightarrow b'}{v b \rightarrow v b'} \end{array}$$

Generally, terms that go wrong are those that “get stuck” during reduction: they are not values yet they don’t reduce.

Example: $1\ 2$ and $\text{add } 1\ (\lambda x.x)$ don’t reduce and are not values.

(See my lecture of 6/2/2020, “Semantics of a functional language”)

Type safety in a reduction semantics

An execution of the program a is viewed as a sequence of reductions:

Termination: $a \rightarrow \dots \rightarrow v \in \text{Val}$

Divergence: $a \rightarrow \dots \rightarrow a' \rightarrow \dots$

Going wrong: $a \rightarrow \dots \rightarrow b \not\rightarrow$ with $b \notin \text{Val}$

Type safety = if $\emptyset \vdash a : \tau$, the “going wrong” case cannot occur.

The standard proof:

- Preservation: if $a \rightarrow a'$ and $\emptyset \vdash a : \tau$ then $\emptyset \vdash a' : \tau$.
- Progress: if $\emptyset \vdash a : \tau$, then $a \in \text{Val}$ or $\exists a'. a \rightarrow a'$.

Run-time checks and “normal” errors

A simple type system cannot rule out all sources of run-time errors, especially

- out-of-bounds array accesses;
- arithmetic errors: divide by zero, overflow.

These errors are detected at run-time (dynamic checking) and reported

- either by aborting the program execution
- or by raising an exception, which can be handled by the program.

In both cases, the program does not go wrong!
It just fails a run-time check.

Reduction semantics with errors

$$\text{div } n \ 0 \rightarrow \text{err} \qquad \frac{n_2 \neq 0 \quad n = n_1/n_2}{\text{div } n_1 \ n_2 \rightarrow n}$$

Error propagation across execution:

$$\text{err } a \rightarrow \text{err}$$

$$v \ \text{err} \rightarrow \text{err}$$

Handling the error:

$$\text{handle } v \ a \rightarrow v$$

$$\text{handle } \text{err } a \rightarrow a$$

$$\frac{a \rightarrow a'}{\text{handle } a \ b \rightarrow \text{handle } a' \ b}$$

Type safety in a reduction semantics with errors

A fourth possible outcome for the execution of a program:

Normal termination: $a \rightarrow \dots \rightarrow v \in \text{Val}$

Termination on error: $a \rightarrow \dots \rightarrow \text{err}$

Divergence: $a \rightarrow \dots \rightarrow a' \rightarrow \dots$

Going wrong: $a \rightarrow \dots \rightarrow b \not\rightarrow$ with $b \notin \text{Val}, b \neq \text{err}$

Same definition for type safety:

if $\emptyset \vdash a : \tau$, the “going wrong” case cannot occur.

Similar proof, taking $\emptyset \vdash \text{err} : \tau$

(err is a term that belongs to all types).

What about dynamically-typed languages?

In a dynamic typing approach, there are no “going wrong” situations at run-time, only normal errors.

We can model this by adding error generation rules:

$$\begin{array}{ll} x \rightarrow \text{err} & \text{add } (\lambda x.a) v \rightarrow \text{err} \\ n v \rightarrow \text{err} & \text{add } n (\lambda x.a) \rightarrow \text{err} \end{array}$$

The “going wrong” case becomes impossible:

$$a \rightarrow \dots \rightarrow b \not\rightarrow \text{ with } b \notin \text{Val}, b \neq \text{err}$$

since every term that is neither a value nor `err` can reduce.
Therefore, dynamic typing is type safe by construction...

From “going wrong” to “undefined behavior”

The standard formalization of “going wrong” suggests an execution that crashes and stops immediately:

$$\mathcal{E}(\eta, a) = \text{wrong} \quad \text{or} \quad a \not\rightarrow, a \notin \text{Val}, a \neq \text{err}.$$

This is not a major security risk!

(No more than stopping on a normal error $a \rightarrow \text{err}$.)

The major risks are executing arbitrary code, producing a wrong result, revealing a secret, etc.

The C and C++ standards use the notion of **undefined behavior** to state that **anything can happen** when the program goes wrong.

Modeling undefined behavior

First try: the problematic terms can reduce to any term a .

$$n \ v \rightarrow a \quad \text{add } (\lambda x.b) \ v \rightarrow a \quad \text{add } n \ (\lambda x.b) \rightarrow a$$

Limitation: some undefined behaviors cannot be expressed by a term of the language.

(For instance: performing I/O in a pure language.)

Inconvenience: it's hard to distinguish reduction sequences $a \rightarrow a_1 \rightarrow \dots \rightarrow a_n \rightarrow \dots$ where “everything is fine” from those that trigger undefined behavior.

Modeling undefined behavior

Alternative: problematic terms reduce to a special `wrong` term, which then can reduce to any term a .

$$n \ v \rightarrow \text{wrong} \quad \text{add } (\lambda x. b) \ v \rightarrow \text{wrong} \quad \text{add } n \ (\lambda x. b) \rightarrow \text{wrong}$$

for all a , $\text{wrong} \rightarrow a$

Reduction sequences $a \rightarrow a_1 \rightarrow \dots \rightarrow a_n \rightarrow \dots$ where “everything is fine” are those that do not contain `wrong`.

The usual proof of type safety can be easily adapted:

- Preservation: if $a \rightarrow a'$ and $\emptyset \vdash a : \tau$ then $a' \neq \text{wrong}$ and $\emptyset \vdash a' : \tau$.
- Progress: if $\emptyset \vdash a : \tau$, then $a \in \text{Val}$ or $\exists a'. a \rightarrow a'$, but $a \neq \text{wrong}$.

Type abstraction

Types for hiding representations

The meaning of a syntactically-valid program in a “type-correct” language should never depend upon the particular representation used to implement its primitive types.

J. C. Reynolds, Towards a theory of type structure, 1974.

The type system of a language distinguishes base types even when they have same machine-level representation:

integer \neq float \neq reference (repr: 64-bit words)

string \neq code of a function (repr: array of bytes)

Functional encapsulation

```
let next =  
  let counter = ref 0 in  
  fun () -> incr counter; !counter
```

By lexical scoping, the `counter` reference is only accessible to the `next` function.

In memory, `next` is represented by a function closure: a pair (code pointer, free variable `counter`).

Without strong typing, anyone could access `counter`:

```
let extracted_counter = snd (next : unit * int ref)
```

The meaning of a syntactically-valid program in a “type-correct” language should never depend upon the particular representation used to implement its primitive types. [...]

The main thesis of Morris (1971) is that this property of representation independence should hold for user-defined types as well as primitive types.

J. C. Reynolds, Towards a theory of type structure, 1974.

Type abstraction: a linguistic mechanism to **hide the concrete representation** of a program-defined data type, forcing users of this type to **go through the operations provided over the type**.

Capabilities as an abstract type

```
module Capa:  
  : sig type t  
    val init: unit -> t  
    val allowed: permission -> t -> bool  
    val drop: permission -> t -> t  
  end  
= struct type t = permission list ... end
```

The signature constraint “hides” the fact that `Capa.t` is implemented as `permission list`.

For the clients of `Capa`, the type `Capa.t` is as “opaque” as `float` or `int → int`.

The only possible values of type `Capa.t` are those obtained by applying `Capa.init` and `Capa.drop`.

Type abstraction in object-oriented languages

In Java, similar guarantees can be achieved by hiding (with the help of visibility modifiers) the internal state and the default constructors of a class.

```
public final class Capa {  
    private T capa;  
    private Capa(T p) { this.capa = p; }  
    public static Capa init() { return new Capa(...); }  
    public bool allowed(int p) { ... }  
    public Capa drop(int p) { ... }  
}
```

Type abstraction and run-time safety

A type system can guarantee run-time safety (in Milner's sense, *well-typed programs do not go wrong*) yet fail to enforce type abstraction. Example:

```
module Capa:  
  : sig type t ... end  
  = struct type t = permission list ... end
```

SML signature constraints have a different meaning than in OCaml. In SML, the type-checker reveals to clients of `Capa` that `Capa.t = permission list`.

The client can, therefore, construct a list `[p1; p2]` of permissions and pass it to any function that expects a `Capa.t`.

This breaks security, but not run-time safety!

Representation independence

(J. C. Reynolds, *Types, abstraction and parametric polymorphism*, 1983)

Respecting type abstraction is not a property of one run of a client of the abstraction: it's a **hyperproperty** of **two runs** of the client, linked with **two different implementations** of the abstraction.

This hyperproperty is called **representation independence**: it must be possible to replace one implementation of an abstract type (e.g. `Capa.t = permission list`) with another implementation, without changing the behaviors of the clients of the abstraction.

Two implementations of the abstraction Capa

```
type permission = P0 | P1 | P2
```

```
module Capa1 = struct
```

```
type t = permission list
```

```
let init () = [P0;P1;P2]
```

```
let allowed = List.mem
```

```
let drop = List.remove
```

```
end
```

```
module Capa2 = struct
```

```
type t = int
```

```
let mask = function
```

```
    P0 -> 1 | P1 -> 2 | P2 -> 4
```

```
let init () = 7
```

```
let allowed p c =
```

```
    c land mask p <> 0
```

```
let drop p c =
```

```
    c land lnot (mask p)
```

```
end
```

A relation between the two implementations

Idea: let's construct a **relation** between the two implementations, telling when a `permission list` and an `int` represent the same abstract set of permissions.

$$\begin{aligned} V(\text{Capa.t}) = \{ (L, n) : \text{permission list} \times \text{int} \mid \\ & (P0 \in L \Leftrightarrow \text{bit}(n, 0) = 1) \\ & \wedge (P1 \in L \Leftrightarrow \text{bit}(n, 1) = 1) \\ & \wedge (P2 \in L \Leftrightarrow \text{bit}(n, 2) = 1) \} \end{aligned}$$

A logical relation

We then extend this relation $V(t)$ between values to all types t

$$V(\text{int}) = \{ (n, n) \mid n \text{ integer} \}$$

$$V(t \rightarrow s) = \{ (\lambda x_1. a_1, \lambda x_2. a_2) \mid \\ \forall (v_1, v_2) \in V(t), (a_1[x_1 \leftarrow v_1], a_2[x_2 \leftarrow v_2]) \in E(s) \}$$

We then extend it from values to terms (computations)

$$E(t) = \{ (a_1, a_2) \mid \forall b_1, a_1 \xrightarrow{*} b_1 \wedge b_1 \text{ irreducible} \Rightarrow \\ \exists b_2, a_2 \xrightarrow{*} b_2 \wedge (b_1, b_2) \in V(t) \}$$

Intuition: if $(a_1, a_2) \in E(t)$, the computation a_1 linked with the first implementation of Capa behaves exactly like the computation a_2 linked with the second implementation.

The fundamental theorem of logical relations

In a well-typed term a , free variables x_i can be interpreted by related values v_i, v'_i , and the two computations we obtain are related.

Theorem (logical relations)

If $x_1 : \tau_1, \dots, x_n : \tau_n \vdash a : \tau$ and $(v_i, v'_i) \in V(\tau_i)$ for every i , then

$$(a\{x_i \leftarrow v_i\}, a\{x_i \leftarrow v'_i\}) \in E(\tau)$$

This result is strictly stronger than type soundness: it shows not only run-time safety, but also representation independence.

Relating the two implementations of `Capa`

We show that the operations of the two implementations are related at their types:

$$\begin{aligned} & ((\text{fun } () \rightarrow [P0;P1;P2]), \text{fun } () \rightarrow 7) \\ & \in V(\text{unit} \rightarrow \text{Capa.t}) \end{aligned}$$
$$\begin{aligned} & (\text{List.mem}, \text{fun } p \ c \rightarrow c \ \text{land} \ \text{mask } p \ <> \ 0) \\ & \in V(\text{permission} \rightarrow \text{Capa.t} \rightarrow \text{bool}) \end{aligned}$$
$$\begin{aligned} & (\text{List.remove}, \text{fun } p \ c \rightarrow c \ \text{land} \ \text{lnot } (\text{mask } p)) \\ & \in V(\text{permission} \rightarrow \text{Capa.t} \rightarrow \text{Capa.t}) \end{aligned}$$

(Just observe that related arguments are mapped to related results.)

Representation independence for Capa

Assume $a : \text{int}$ under the typing hypotheses

$\text{Capa.init} : \text{unit} \rightarrow \text{Capa.t}$

$\text{Capa.allowed} : \text{permission} \rightarrow \text{Capa.t} \rightarrow \text{bool}$

$\text{Capa.remove} : \text{permission} \rightarrow \text{Capa.t} \rightarrow \text{Capa.t}$

Let a_1, a_2 be the programs obtained by linking a with one of the implementations of Capa :

$$a_1 = a\{\text{Capa} \leftarrow \text{Capa1}\} \quad a_2 = a\{\text{Capa} \leftarrow \text{Capa2}\}$$

Then, $(a_1, a_2) \in E(\text{int})$.

This means that both programs evaluate to the same integer.

Static typing of resources

Explicit memory management

Dynamic memory allocation + **explicit deallocation** under the programmer's control.

Example: `malloc` and `free` in C, `new` and `delete` in C++.

```
p = malloc(10);  
/* use p */;  
free(p);
```

A source of many programming errors!

Memory leak:

```
p = malloc(10);  
if ... else return;  
free(p);
```

Use after free:

```
p = malloc(10);  
...  
free(p);  
/* use p */
```

Double free:


```
p = malloc(10);  
...  
free(p);  
...  
free(p);
```

An attack on use-after-free



Allocate a large array p .


An attack on use-after-free



Allocate a large array p .

Free it immediately.

An attack on use-after-free




Allocate a large array p .

Free it immediately.

Wait for the memory area to be reused for other allocations
(of sensitive data).

An attack on use-after-free



Allocate a large array p .

Free it immediately.

Wait for the memory area to be reused for other allocations (of sensitive data).

Read or modify sensitive data from pointer p .

Note: this invalidates run-time safety, even if the language is strongly typed!

Automatic memory management

No explicit deallocation of memory by the program, but **automatic reclamation by the run-time system** of the memory blocks that are no longer reachable.

(Reference counting, garbage collection, etc.)

Ex: Lisp, functional languages, scripting languages, Java, Go, ...

For a long time, automatic memory management was believed to be necessary for type safety.

Limitations:

- Not always applicable
(e.g. in an OS or within the implementation of a memory manager).
- Other kinds of **resources** still need manual management.

Example of resources: file descriptors

A familiar API to read files:

```
open : string → file
input_line : file → string
close : file → unit
```

A typical use:

```
let f = open "foo" in
let l = input_line f in
close f; l
```

Incorrect handling of file descriptors

A possible leak of a file descriptor:

```
let f = open "foo" in let l = input_line f in close f; l
```

If file `foo` is empty, `input_line f` raises an exception and `f` is not closed.

A read after close:

```
let f = open "foo" in ... close f; ...; input_line f
```

A double close:

```
let f = open "foo" in ... close f; ...; close f
```

Typically, these errors are detected at run-time.

Aliasing and resource sharing

```
let interleave f1 f2 =  
  ... input_line f1 ... input_line f2 ...;  
  close f1; close f2
```

If `f1` and `f2` are **aliases** on the same descriptor, we have a double close.

```
let interleave flist =  
  ... List.map input_line flist ...  
  List.iter close flist  
in  
  let f = open "foo" in  
  let l1 = [f; open "gee"] and l2 = [f; open "buz"] in  
  interleave l1; interleave l2
```

`f` is **shared** between the two lists `l1` and `l2`, causing a read-after-close in `interleave l2`.

Controlling resources by static typing

An idea that appeared in the pure functional language community, where (morally) we never modify a resource; instead, we return a modified resource.

```
open : string → file
input_line : file → string * file
close : file → unit
```

A new problem appears: we must not use a file value twice!

```
let f1 = open "foo" in
let (l1, f2) =
  input_line f1 in ✓
let (l2, f3) =
  input_line f2 in ✓
...
```

```
let f1 = open "foo" in
let (l1, f2) =
  input_line f1 in ✓
let (l2, f3) =
  input_line f1 in ✗
...
```

Uniqueness types in the Clean language

The type `unique τ` of values of type τ that are reachable via one reference only, and that can therefore be implemented using in-place modification.

```
open : string → unique file
input_line : unique file → string * unique file
close : unique file → unit
```

Prevents incorrect reuse of values:

```
let f1 = open "foo" in
let (l1, f2) = input_line f1 in
let (l2, f3) = input_line f1 in
```

Two uses of `f1 : unique file`, rejected by type-checking.

Leaks are not prevented: we can still forget to close `f3`.

(Inspired by Girard's linear logic.)

A type $\sigma \multimap \tau$ of functions that use their σ argument exactly once.

```
open :  $\forall \alpha. \text{string} \rightarrow (\text{file} \multimap \alpha) \rightarrow \alpha$   
input_line :  $\text{file} \multimap \text{string} * \text{file}$   
close :  $\text{file} \multimap \text{unit}$ 
```

Prohibits multiple uses of a `file` value and forces us to call `close` at the end.

(This would not be the case for `open : string \rightarrow file`.)

Tracing resource ownership

Another approach to resource management, originating in object-oriented languages, popularized by the Rust language.

```
open: string → file  
close: file → unit
```

In the simplest case, a resource is owned by a variable and is automatically freed at the end of the variable scope.

```
begin let f = open "foo" in  
...  
(* implicit call to close f *)  
end
```

Tracing resource ownership

Another approach to resource management, originating in object-oriented languages, popularized by the Rust language.

```
open: string → file
close: file → unit
```

The resource can also be explicitly transferred to a function, in which case it is no longer owned by the variable.

```
begin let f = open "foo" in
...
close f
(* no implicit call to close f *)
end
```

Tracing resource ownership

Another approach to resource management, originating in object-oriented languages, popularized by the Rust language.

```
open: string → file  
close: file → unit
```

After transfer, we cannot use the resource any longer, nor transfer it again.

```
let f = open "foo" in  
...  
close f;  
close f ❌
```

Tracing resource ownership

Another approach to resource management, originating in object-oriented languages, popularized by the Rust language.

```
open: string → file  
close: file → unit
```

Taking an alias on a resource is treated like a transfer.

```
let f = open "foo" in  
let g = f in  
...  
close f; ✗  
close g  
end
```

Borrowing a resource

```
open: string → file
close: file → unit
input_line: &mut file → string
position: & file → int
```

A **borrow** gives temporary right to use the resource.

```
let f = open "foo" in
let l = input_line (&mut f) in
close f; l
```

The resource `f` is “lent” to `input_line`, then recovered when this function returns.

Borrowing a resource

```
open: string → file
close: file → unit
input_line: &mut file → string
position: & file → int
```

During a mutable borrow, the original owner cannot perform any action on the resource.

```
let f = open "foo" in
let b = &mut f in
close f; ✗
input_line b
```

```
let f = open "foo" in
let b = &mut f in
let l = input_line (&mut f) in ✗
input_line b
```

Borrowing a resource

```
open: string → file
close: file → unit
input_line: &mut file → string
position: & file → int
```

Several immutable borrows can be ongoing at the same time.
(The “mutable XOR shared” policy.)

```
let f = open "foo" in
let b1 = & f in let b2 = & f in
position b1 + position b2
```


Application: zero-copy message passing

```
send: buffer → unit  
receive: unit → buffer
```

Passing messages between two threads running concurrently:

```
let b = new_buffer() in    ||    let b = receive() in  
fill(&mut b);              ||    if check(& b)  
send(b)                    ||    then use(&mut b)  
                            ||    else error()
```

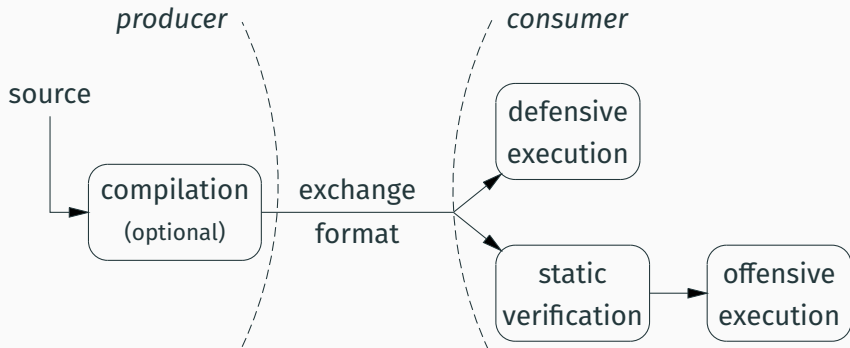
After `send(b)`, the left thread cannot operate on `b`

→ no Time Of Check To Time Of Use attack

where `b` would change between `check(& b)` and `use(&mut b)`.

Typing and verification of mobile code

Mobile code



Mobile code formats: source (JavaScript), intermediate (JVM bytecode), native (x86 machine code).

What static verifications can we perform on machine codes, either virtual machine(JVM) or hardware processors (x86)?

Java bytecode verification

A static analysis on the JVM bytecode intermediate representation that establishes many useful properties before execution:

- **Code is well formed.**
(E.g. no branch in the middle of another method.)
- **Instructions receive arguments of the expected types.**
(E.g. `getField C.f` receives an object of class `C` or a sub-class.)
- **The expression stack does not overflow.**
(Within one method; dynamic check at method calls.)
- **Local variables (registers) are initialized before use.**
(No access to values that remain from an earlier call.)
- **Objects (class instances) are initialized before use.**
(I.e. `new C`, then call to a constructor, then use.)
- **Visibility modifiers are respected.**
(E.g. no access to a `private` member outside of the defining class.)

A number of verifications that are crucial for run-time safety and for security are still performed at run-time:

- bounds checks for array accesses;
- checks for null references;
- conversion to a sub-class (down-casting);
- typing stores into arrays of objects;
- stack inspection by the `SecurityManager`.

Verifying branchless code

Executing the JVM code by a **defensive abstract machine** that uses types in place of values.

- The machine maintains a stack of types and a set of registers containing types.
- For each instruction, it checks the types of the arguments, computes the type of the result, and updates the type of the destination.

Example:

```
class C {  
    int x;  
    void move(int delta) {  
        int oldx = x; x += delta; D.draw(oldx, x);  
    }  
}
```

```
ALOAD 0  
  
DUP  
  
GETFIELD C.x : int  
  
DUP  
  
ISTORE 2  
  
ILOAD 1  
  
IADD  
  
SETFIELD C.x : int  
  
ILOAD 2  
  
ALOAD 0  
  
GETFIELD C.x : int  
  
INVOKESTATIC D.draw : void(int,int)  
  
RETURN
```

r0: C, r1: int, r2: T []

ALOAD 0

DUP

GETFIELD C.x : int

DUP

ISTORE 2

ILOAD 1

IADD

SETFIELD C.x : int

ILOAD 2

ALOAD 0

GETFIELD C.x : int

INVOKESTATIC D.draw : void(int,int)

RETURN


```
                                r0: C, r1: int, r2: T   [ ]
ALOAD 0
                                r0: C, r1: int, r2: T   [ C ]
DUP
GETFIELD C.x : int
DUP
ISTORE 2
ILOAD 1
IADD
SETFIELD C.x : int
ILOAD 2
ALOAD 0
GETFIELD C.x : int
INVOKESTATIC D.draw : void(int,int)
RETURN
```

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
DUP		
ISTORE 2		
ILOAD 1		
IADD		
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
ISTORE 2		
ILOAD 1		
IADD		
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
	r0: C, r1: int, r2: T	[C; int; int]
ISTORE 2		
ILOAD 1		
IADD		
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		


	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
	r0: C, r1: int, r2: T	[C; int; int]
ISTORE 2		
	r0: C, r1: int, r2: int	[C; int]
ILOAD 1		
IADD		
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
	r0: C, r1: int, r2: T	[C; int; int]
ISTORE 2		
	r0: C, r1: int, r2: int	[C; int]
ILOAD 1		
	r0: C, r1: int, r2: int	[C; int; int]
IADD		
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
	r0: C, r1: int, r2: T	[C; int; int]
ISTORE 2		
	r0: C, r1: int, r2: int	[C; int]
ILOAD 1		
	r0: C, r1: int, r2: int	[C; int; int]
IADD		
	r0: C, r1: int, r2: int	[C; int]
SETFIELD C.x : int		
ILOAD 2		
ALOAD 0		
GETFIELD C.x : int		
INVOKESTATIC D.draw : void(int,int)		
RETURN		

	r0: C, r1: int, r2: T	[]
ALOAD 0		
	r0: C, r1: int, r2: T	[C]
DUP		
	r0: C, r1: int, r2: T	[C; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: T	[C; int]
DUP		
	r0: C, r1: int, r2: T	[C; int; int]
ISTORE 2		
	r0: C, r1: int, r2: int	[C; int]
ILOAD 1		
	r0: C, r1: int, r2: int	[C; int; int]
IADD		
	r0: C, r1: int, r2: int	[C; int]
SETFIELD C.x : int		
	r0: C, r1: int, r2: int	[]
ILOAD 2		
	r0: C, r1: int, r2: int	[int]
ALOAD 0		
	r0: C, r1: int, r2: int	[int; C]
GETFIELD C.x : int		
	r0: C, r1: int, r2: int	[int; int]
INVOKESTATIC D.draw : void(int,int)		
	r0: C, r1: int, r2: int	[]
RETURN		

Verifying code containing branches



A classic **dataflow analysis**:

Fork points:

propagate types to
all successors.

Join points:

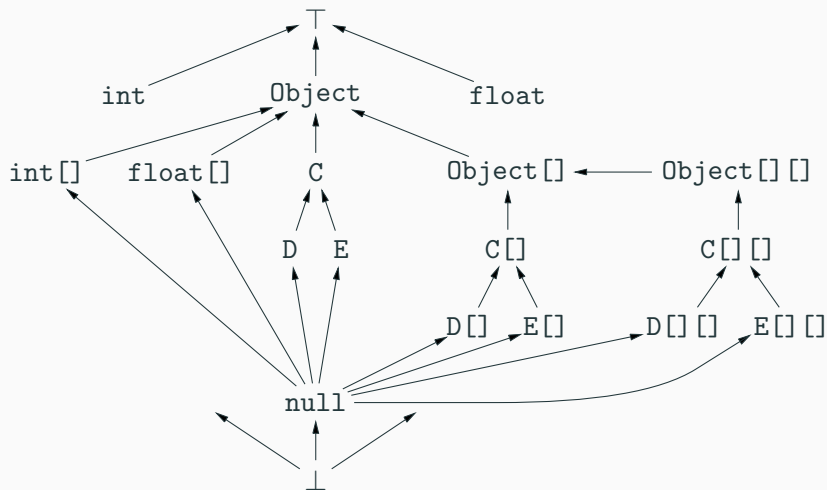
take the least upper bound of
the types of all predecessors.

Iterate the analysis,

until a fixed point is reached.

(See lecture of 19/12/2019.)

The lattice of JVM types (simplified)



Several JVM features complicate bytecode verification beyond a classic dataflow analysis:

- **Interfaces:**
the subtype relation is not a semi-lattice.
- **The object initialization protocol:**
requires a bit of must-alias analysis.
- **Subroutines:**
a code sharing mechanism, no longer in use, that required a polyvariant analysis.


Interfaces

```
interface I { ... }
```

```
interface J { ... }
```

```
class C1 implements I, J { ... }
```

```
class C2 implements I, J { ... }
```



A class can be subtype of several interfaces.

Thus, the subtyping order is not a semi-lattice:

C1 and C2 have two incomparable super-types, I and J.


Interfaces

```
interface I { ... }
```

```
interface J { ... }
```

```
class C1 implements I, J { ... }
```

```
class C2 implements I, J { ... }
```



Dedekind-MacNeille completion: adding points to recover a semi-lattice.

Here, the pseudo-class `IandJ` was added as l.u.b. of `C1` and `C2`.


Interfaces

```
interface I { ... }
```

```
interface J { ... }
```

```
class C1 implements I, J { ... }
```

```
class C2 implements I, J { ... }
```



Java's original solution: the bytecode verifier ignores interfaces, treating them all like `Object`.

A run-time check is performed by the `invokeinterface I.m`, since the argument cannot be statically guaranteed to implement interface `I`.

A more modern approach: verification using certificates

(E. Rose, *Lightweight Bytecode Verification*, 2003. The KVM. The JVM since Java 7.)

The Java compiler annotates the produced JVM bytecode with **stack maps**, i.e. types for the stack and the registers, at some points in the bytecode:

- at the beginning of each basic block (Java 7);
- at each instruction where the types “before” differ from the types “after” the preceding instruction (E. Rose).

Type checking can then be performed in a single linear pass, without fixed-point iteration, and without computing least upper bounds.

(G. Necula, P. Lee, *et al*, 1996-2000)

A general, ambitious approach to the security of mobile code:

- Much freedom to choose the language of the mobile code, all the way to machine code (x86 or other) produced by an optimizing compiler, or hand-written.
- Much freedom to choose a security policy, from type safety to triples $\{ P \} c \{ Q \}$ in a program logic.
- Verifying the code against the policy can be expensive, even undecidable, and involve automated theorem proving.

Proof-Carrying Code

Core idea: separate the verification of the code in two phases:

1. Certification (on the code producer side):
production of a “proof term” or “certificate”
2. Validation (on the code consumer side):
checking consistency between certificate, code, and
expected property.

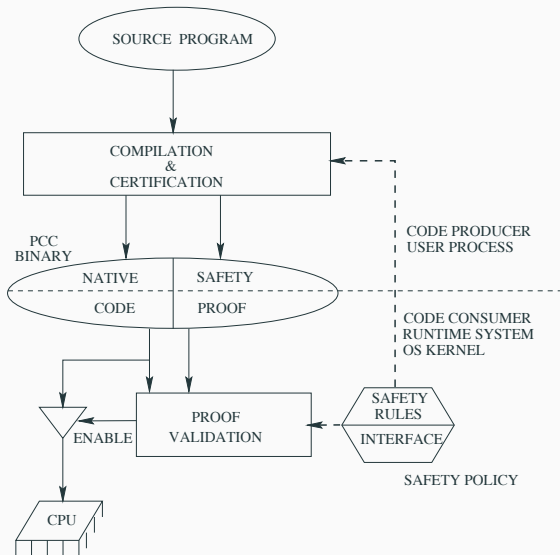
Example: Java bytecode verification.

- Certification: producing stack maps for all basic blocks.
- Validation: type-checking every basic block.

Example: proving a theorem $P : Prop$ in Coq.

- Certification: construction of a proof term $p : P$.
- Validation: type-checking to verify that $p : P$.

The PCC architecture



A fragment of a security policy

(G. Necula, *Compiling with Proofs*, 1998.)

Expressions: $e ::= x \mid e_1 + e_2 \mid e_1 \& e_2 \mid e_1 \mid e_2 \mid \text{sel}(m, e)$

Memory: $m ::= x \mid \text{upd}(m, e_1, e_2)$

Types: $\tau ::= \text{int} \mid \text{bool} \mid \text{array}(\tau, e)$

Predicates: $P ::= P_1 \wedge P_2 \mid P_1 \supset P_2 \mid \forall x. P_x$
 $\mid e_1 \geq 0 \mid e : \tau \mid \text{saferd}(e)$

Rules:

$$\frac{e_1 : \text{bool} \quad e_2 : \text{bool}}{e_1 \& e_2 : \text{bool}} \quad \frac{e_1 : \text{bool} \quad e_2 : \text{bool}}{e_1 \mid e_2 : \text{bool}} \quad \frac{}{\text{sizeof}(\text{bool}) = 1}$$

$$\frac{a : \text{array}(\tau, \text{len}) \quad i \geq 0 \quad i < \text{len} * \text{sizeof}(\tau)}{\text{saferd}(a + i)}$$

$$\frac{a : \text{array}(\tau, \text{len}) \quad \text{sizeof}(\tau) = 1 \quad i \geq 0 \quad i < \text{len}}{\text{sel}(m, a + i) : \tau}$$

Array accesses are within bounds.

The representation of the `bool` type is kept abstract.

Production of optimized machine code

```
bool main(bool A[], bool B) {  
    int I;  
    bool R = B;  
    for(I = 0; I < length(A); I++)  
        R = R && A[i];  
    return R;  
}
```

```
    rI = 0  
    rR = rB  
L0:   INV rI ≥ 0 ∧ rI : int ∧ rR : bool  
    if rI ≥ rL goto Lend  
    rT = *(rA + rI)  
    rR = rR & rT  
    goto L0  
Lend: return rR
```

The compiler generated no run-time bounds check because it detected that the access $A[i]$ is always within A 's bounds.

The compiler annotated the generated code with a loop invariant.

Production of the verification condition

Using a strongest postcondition calculus, with the specification

$$\{r_B : \text{bool} \wedge r_A : \text{array}(\text{bool}, r_L)\} \text{C} \{r_R : \text{bool}\}$$

```
1  $\forall r_A. \forall r_B. \forall r_L. \forall m.$   
2  $r_B : \text{bool} \wedge r_A : \text{array}(\text{bool}, r_L) \supset$   
3  $(0 \geq 0 \wedge 0 : \text{int} \wedge r_B : \text{bool}) \wedge$   
4  $\forall r_I. \forall r_R.$   
5  $r_I \geq 0 \wedge r_I : \text{int} \wedge r_R : \text{bool} \supset$   
6  $(r_I < r_L \supset r_I + 1 : \text{int} \wedge r_R \& \text{sel}(m, r_A + r_I) : \text{bool} \wedge$   
7  $\text{saferd}(r_A + r_I)) \wedge$   
8  $(r_I \geq r_L \supset r_R : \text{bool})$ 
```

Representing and verifying the proof in LF

The *Logical Framework*: a dependently-typed lambda-calculus, able to express propositions and proof terms.

```
exp      : Type
tp       : Type
pred     : Type
pf       : pred → Type

true     : pred
and      : pred → pred → pred
imp      : pred → pred → pred
int      : tp
array    : tp → exp → tp
of       : exp → tp → pred
saferd   : exp → exp → pred

truei    : pf true
andi     :  $\Pi P:\text{pred}.\Pi R:\text{pred}.\text{pf } P \rightarrow \text{pf } R \rightarrow \text{pf } (\text{and } P R)$ 
andel    :  $\Pi P:\text{pred}.\Pi R:\text{pred}.\text{pf } (\text{and } P R) \rightarrow \text{pf } P$ 
szbool   : pf (= (sizeof bool) 1)
rdarray  :  $\Pi M:\text{exp}.\Pi A:\text{exp}.\Pi I:\text{exp}.\Pi L:\text{exp}.\Pi T:\text{tp}.$ 
          pf (of A (array T L)) →
          pf (= (sizeof T) 1) →
          pf ( $\geq I 0$ ) →
          pf ( $< I L$ ) →
          pf (of (sel M (plus A I)) T)
```

Verifying that c is a valid proof of proposition P is easy:
it suffices to check that $c : \text{pf } P$.

The Touchstone compiler

(Colby et al, 2000)

Compilation Java bytecode → optimized x86 code,
producing certificates of type safety.

Native code injection for network packet filtering

(Necula & Lee, 1996)

Like BPF and eBPF, but the code injected in the kernel is native code, and verifying the certificate is simpler than the safety analysis done by eBPF.

Certificates are huge

- Much redundancy in LF terms, can be improved by using implicit arguments.
- Alternate approach for validation: nondeterministic proof search guided by an “oracle”, which is the certificate.

The security policy and the v.c.gen. are part of the trusted computing base


- They must be verified independently.
- *Foundational Proof-Carrying Code*: (Appel et al, 1999-2005)
the typing rules and the program logic are derived from the operational semantics of the machine code.

Summary

What does typing contribute to security?

Strong typing (static or dynamic) provides basic guarantees that are necessary for software security: integrity of executions, data structures, and memory.

For instance, these guarantees would have prevented up to 70% of the serious bugs in the Chrome browser.



What does typing contribute to security?

Strong typing (static or dynamic) provides basic guarantees that are necessary for software security: integrity of executions, data structures, and memory.

Some type systems provide additional guarantees, such as:

- Type abstraction and representation independence.
- Control of the ownership and proper usage of resources.
- Control of information flow; non-interference (lecture #2).

What does typing contribute to security?

Strong typing (static or dynamic) provides basic guarantees that are necessary for software security: integrity of executions, data structures, and memory.

Some type systems provide additional guarantees, such as:

- Type abstraction and representation independence.
- Control of the ownership and proper usage of resources.
- Control of information flow; non-interference (lecture #2).

Difficulties in combining these type-based approaches and to put them into practice.

What does typing contribute to security?

Strong typing (static or dynamic) provides basic guarantees that are necessary for software security: integrity of executions, data structures, and memory.

Some type systems provide additional guarantees, such as:

- Type abstraction and representation independence.
- Control of the ownership and proper usage of resources.
- Control of information flow; non-interference (lecture #2).

Difficulties in combining these type-based approaches and to put them into practice.

An ongoing transition from typing to program proof, already anticipated in Proof Carrying Code.