

Why semantics matter

Richard Bonichon

20210105

Outline

Programming is hard

Overview of a compiler

Compilers are complex

Language oddities zoo

Programming is hard

90%

of programmers
make some kind of errors when
coding **binary search**.

The div of death

```
1 #include <stdio.h>
2
3 int div(int a) {
4     return a / a;
5 }
6
7 int main () {
8     printf("d5 = %d\n", div(5));
9     printf("d0 = %d\n", div(0));
10    return 0;
11 }
```

The answer(s)

-O	gcc 8.3.0	clang 7.1.0
0	d5 = 1 d0 = 1	core dumped
1	d5 = 1 d0 = 1	d5 = 1 d0 = 1
2	d5 = 1 d0 = 1	d5 = 1 d0 = 1
3	d5 = 1 d0 = 1	d5 = 1 d0 = 1

What is printed ?

```
1  #include "stdio.h"
2
3  long foo(int *x, long *y) {
4      *x = 0;
5      *y = 1;
6      return *x;
7  }
8
9  int main(void) {
10     long l;
11     printf("%ld\n", foo((int *) &l, &l));
12     return 0;
13 }
```

Answer(s) on x86

-O	gcc 8.3.0	clang 7.1.0
0	1	1
1	1	0
2	0	0
3	0	0

Overview of a compiler

What is a compiler ?

Definition

A **compiler** is computer software that transforms computer code written in one programming language (the **source language**) into another programming language (**the target language**).

A usual reason for wanting to transform source code is to create an executable program.

```
1|gcc -O2 -S -c simple.c
```

Compilation

```
int foo(int i, int j, int n) {  
    int l, k = 0;  
    for (l = i; l < n; l++)  
        k += i * j;  
    return k;  
}
```

\$ gcc -O2 -S -c simple.c



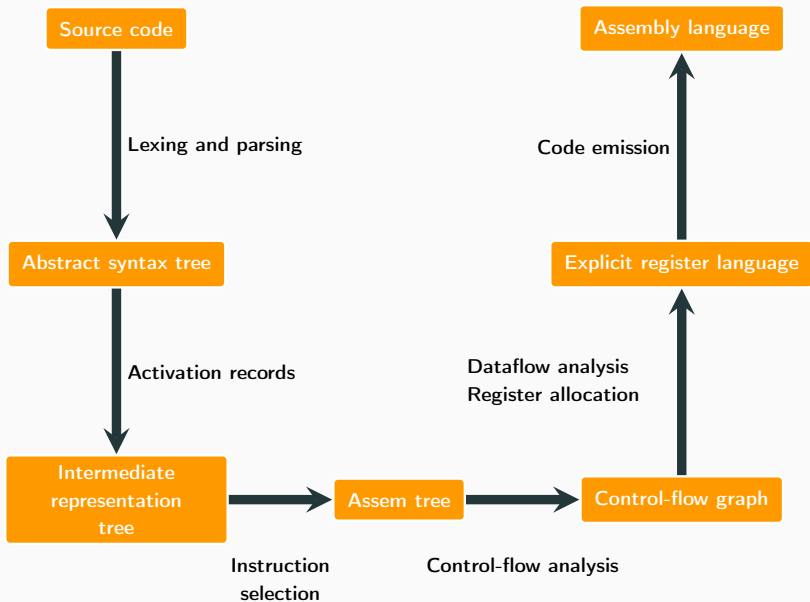
```
foo:  
.LFBO:  
    .cfi_startproc  
    cmpl    %edx, %edi  
    jge    .L3  
    imull   %edi, %esi  
    leal   -1(%rdx), %eax  
    subl   %edi, %eax  
    imull   %esi, %eax  
    addl   %esi, %eax  
    ret  
    .p2align 4,,10  
    .p2align 3  
.L3:  
    xorl   %eax, %eax  
    ret
```

Fair warning

A sufficiently advanced compiler is indistinguishable from an adversary.

– John Regehr

Architecture of a modern compiler



Compilers **must** preserve the semantics of the original program through its many passes.

Definition

- Semantics detail the **meaning** of the program (its statements, expressions, ...)
- Formal semantics interpret programs using mathematics

Understanding a programming language

- what we can trust as regular programmers
- what we need to give as compiler programmers

Tool for designing languages

Fundamentals to show/prove properties of programs

Different types of semantics

Operational semantics

- What the program computes
- Concrete

Denotational semantics

- What the program computes
- Abstract

Axiomatic semantics

- Properties of programs

Example of operational semantics

Semantic rules for a simple imperative language without loops

$$\frac{}{\langle x := a, s \rangle \rightarrow s[x \mapsto \mathcal{A}[[a]]s]} \text{Assigns}$$

$$\frac{}{\langle \text{skip}, s \rangle \rightarrow s} \text{Skip}$$

$$\frac{\langle S_1, s \rangle \rightarrow s' \quad \langle S_2, s' \rangle \rightarrow s''}{\langle S_1; S_2, s \rangle \rightarrow s''} \text{Seq}$$

$$\frac{\langle S_1, s \rangle \rightarrow s' \quad \llbracket b \rrbracket s = \top}{\langle \text{if } b \text{ then } S_1 \text{ else } S_2, s \rangle \rightarrow s'} \text{If}_{\top}$$

$$\frac{\langle S_2, s \rangle \rightarrow s' \quad \llbracket b \rrbracket s = \perp}{\langle \text{if } b \text{ then } S_1 \text{ else } S_2, s \rangle \rightarrow s'} \text{If}_{\perp}$$

Goal

Break the input into lexical unit (**tokens**)

"Does the teacher like compilation ?"

⇒

"Does", "the", "teacher", "like", "compilation", "?"

Goal

Check the structure of sentences (i.e. the **grammar**)

A question of the form

Auxiliary/modal subject (main verb) (direct object) (question mark)

is grammatically valid.

Keywords

Lexing

- Regular expressions
- NFA
- DFA
- Minimization

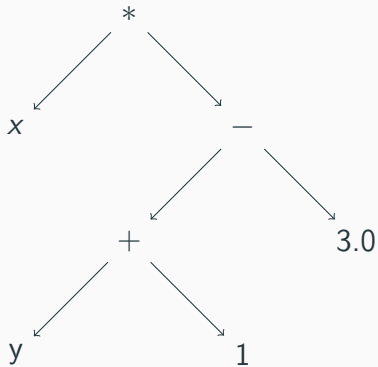
Parsing

- BNF
- LL(k)
- LR(k)

Lexing and parsing transform a **concrete** syntax tree into an **abstract** syntax tree.

Abstract Syntax Tree : $x * ((y + 1) - 3)$

$x * ((y + 1) - 3)$



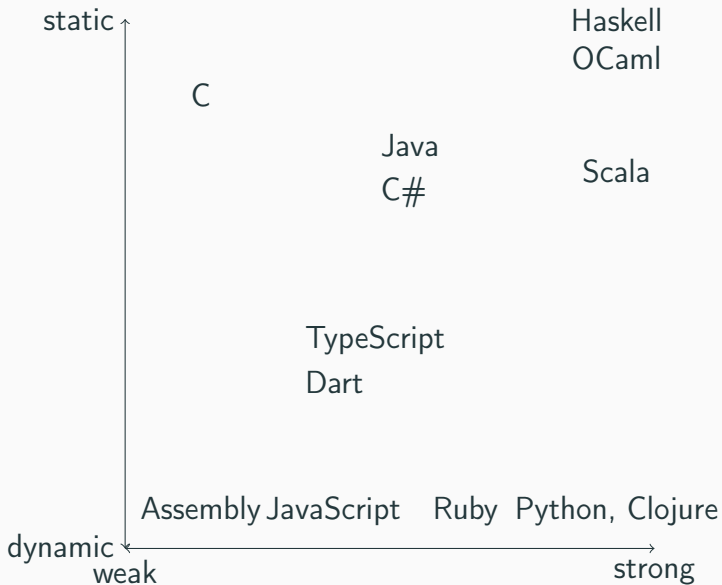
Definition (Typing)

Typing consists in attributing **types** to the data of the program

What for ?

Guarantee that programs make sense, i.e. are valid programs.

Typing systems landscape (Odersky)



Example

$$\frac{\Gamma \vdash b : \text{bool} \quad \Gamma \vdash E_1 : T \quad \Gamma \vdash E_2 : T}{\Gamma \vdash \text{if } b \text{ then } E_1 \text{ else } E_2 : T} \text{If}$$

$$\frac{\Gamma \vdash x : T \quad \Gamma \vdash e : T}{\Gamma \vdash x := e : \text{unit}} \text{Assigns}$$

Intermediate representation

What is an IR ?

Intermediate representation (IR) is the data structure or code used internally by a compiler to represent source code, usually for further processing (optimization, translation)

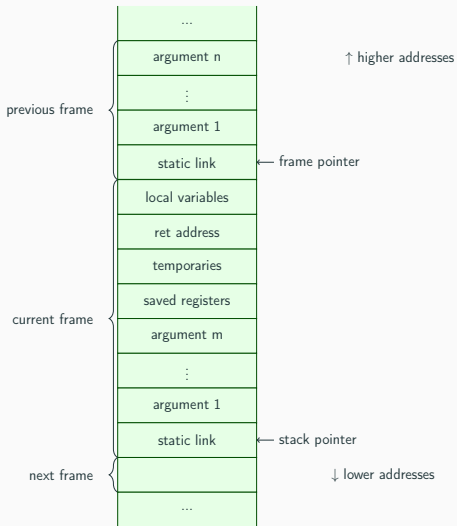
A good IR must be:

- accurate (no loss of information)
- independent of source/target languages.

Examples

- LLVM
- Gimple

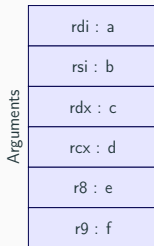
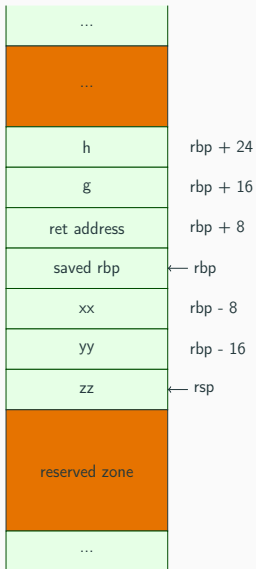
Stack frame allocation



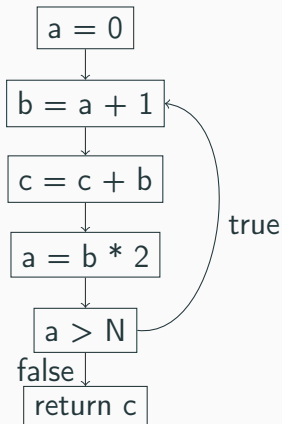
Example

```
1 long myfunc(long a, long b, long c, long d,  
2             long e, long f, long g, long h)  
3 {  
4     long xx = a * b * c * d * e * f * g * h;  
5     long yy = a + b + c + d + e + f + g + h;  
6     long zz = utilfunc(xx, yy, xx % yy);  
7     return zz + 20;  
8 }
```

x86-64 stack on calling myfunc



Control-Flow Graph (CFG)



CFG construction in a nutshell

Howto

- Every node has one statement;
- A directed edge connects two nodes N and M whenever M can be executed right after N in the program

Remarks

In order to know if one statement can follow another, one needs **precise semantics**!

CFGs can be constructed directly from the AST or after it: it is a basic data structure of compilation or static analysis.

CFGs provide a means to compute **reachability** of a given program part. An unreachable code in the CFG:

- will never ever be executed and
- can safely be removed from the program at compile time (this is **dead code**).

Optimizations are done on the CFG through data-flow analyses.

Common subexpression elimination

Definition

Given a statement

- $s : t \leftarrow x \odot y$,

where the expression $x \odot y$ is available at s ,

the computation within s can be eliminated.

Example CSE

Before

```
1 a = b * c + g;  
2 d = b * c * e;
```

After

```
1 tmp = b * c;  
2 a = tmp + g;  
3 d = tmp * e;
```

Constant/copy propagation

Definition

Suppose we have a statement $s_1 : x \leftarrow t$,

where t is either a **constant**, or a **simple variable**.

And another : $s_2 : y \leftarrow x \text{ bop } z$.

x is constant in s_2 if:

- s_1 reaches s_2 **and**
- no other definition of x reaches s_2

In this case : $s_2 : y \leftarrow t \text{ bop } z$

Example

Before

```
1 t = 12;  
2 x = 4;  
3 y = t;  
4 z = x * y - t;
```

After

```
1 t = 12;  
2 x = 4;  
3 y = t;  
4 z = 36;
```

Dead code elimination

Definition

If there is a quadruple

- $s : a \leftarrow b \odot c$; or
- $s : a \leftarrow M[x]$,

such that a is **not live-out** of s ,

then the quadruple can be **deleted**.

Example

Before

```
1 t = 12;  
2 x = 4;  
3 y = t;  
4 z = 36;
```

After

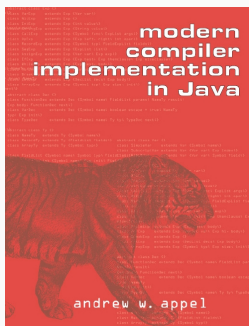
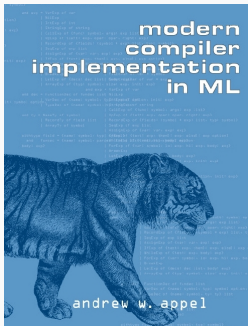
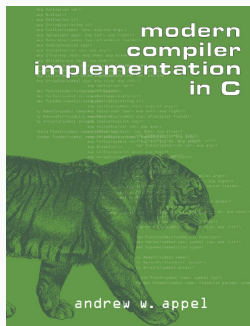
```
1 z = 36;
```

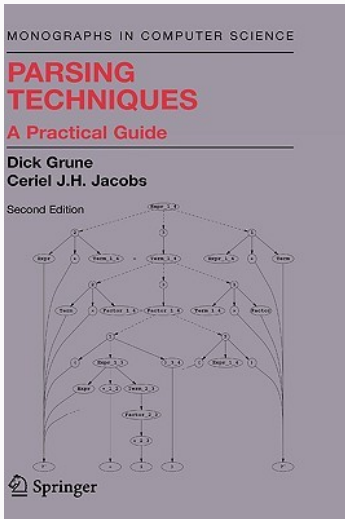
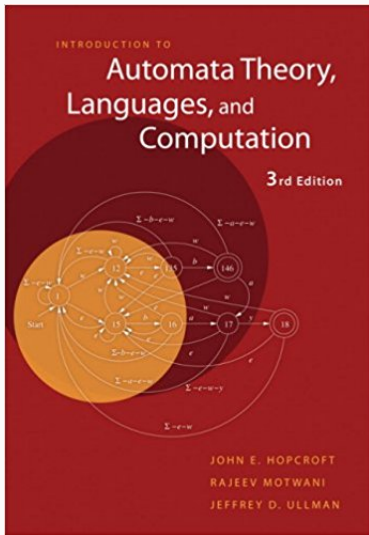

Register allocation

This phase needs to assign:

- • the many temporaries to a small number of concrete machine registers;
- • — where possible — the source and destination of a MOVE to the same register so that the MOVE can be deleted.

More on compilation





Compilers are complex

Source of errors

Coding in C (for example) exposes the programmer to several difficulties

1. Tricky semantics
2. Unforeseen optimizations
3. Undefined behaviors (might be seen as tricky semantics), due to the fact that it is an **unsafe** language

Only 3 is specific to C ...

In particular

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     unsigned char a = 0xff;
6     char b = 0xff;
7     int c = a == b; // true, or false?
8     printf("c = %d\n", c);
9 }
```

Division by zero

```
1 /* Linux kernel : lib/mpi/mpi-pow.c */  
2  
3 if (!msize)  
4     msize = 1 / msize; /* provoke a signal */
```

Oversized shift (Fix for CVE-2009-4307)

```
1  /* Linux kernel: fs/ext4/super.c */
2
3  groups_per_flex = 1 << sbi->s_log_groups_per_flex;
4  /* There are some situations, after shift the value of
5  'groups_per_flex' can become zero and division with 0
6  result in fixpoint divide exception.
7  */
8  if (groups_per_flex == 0) return 1;
9
10 flex_group_count = ... / groups_per_flex;
```


Silent breakage (from Regehr)

```
1 #include <limits.h>
2 #include <stdio.h>
3
4 int foo(int x) {
5     return (x + 1) > x;
6 }
7
8 int main(void) {
9     printf("%d\n", (INT_MAX + 1) > INT_MAX);
10    printf("%d\n", foo(INT_MAX));
11    return 0;
12 }
```

`INT_MAX + 1` is both larger and not larger than `INT_MAX`.

Mixing signed/unsigned

```
1 #include <stdio.h>
2
3 int main (void)
4 {
5     long a = -1;
6     unsigned b = 1;
7     printf ("%d\n", a > b);
8     return 0;
9 }
```

Why, oh why ?

CPUs are typically fastest on integers at their native size.

On x86, 32-bit arithmetic can be twice as fast as 16-bit one.

C is a language focused on performance, so it will do the integer promotion to make the program as fast as possible.

In a nutshell

Keep integer promotion rules in mind to avoid integer overflow vulnerability issues.

Undefined behaviors

Some C operations are left implementation-defined but other are **undefined** in the Standard.

C compilers trust the programmer not to submit code with undefined behaviors.

They optimize code under such assumptions.

Permissible undefined behavior ranges from ignoring the situation completely, with unpredictable results, to having demons fly out of your nose.

But it works on my computer !

Somebody once told me that in basketball you can't hold the ball and run.

I got a basketball and tried it and it worked just fine.

He obviously didn't understand basketball.

– Roger Miller

Why is it good and bad ?

Good

- Makes compiler's job easier

For example, loop optimizations do not have to worry about signed integers overflows — it is undefined behavior.

Bad

- 191 kinds of undefined behaviors in C99

Security problems

```
1 void process_something(int size)
2 {
3     // Catch integer overflow.
4     if (size > size + 1) abort();
5     ... // Error checking from this code elided.
6
7     char *string = malloc(size + 1);
8     read(fd, string, size);
9     string[size] = 0;
10    do_something(string);
11    free(string);
12 }
```

Optimization is hard

```
1 void contains_null_check(int *p)
2 {
3     int dead = *p;
4     if (p == 0)
5         return;
6     *p = 4;
7 }
```


Unwanted dead code elimination

```
1 void check_password(char *pwd);  
2  
3 void get_password(void)  
4 {  
5     char pwd[64];  
6     if (retrieve_password(pwd, sizeof(pwd))) {  
7         check_password(pwd);  
8     }  
9     memset(pwd, 0, sizeof(pwd));  
10 }
```

What is returned ?

```
1  #include <iostream>
2  #include <complex>
3  using namespace std;
4
5  int main() {
6      complex<int> delta;
7      complex<int> mc[4] = {0};
8      int di;
9
10     for(di = 0; di < 4; di++, delta = mc[di]) {
11         cout << "di:" << di << endl;
12         cout << "delta: " << delta << endl;
13     }
14     cout << "mc[di]:" << mc[di] << endl;
15     return 0;
16 }
```

Take away

At low-level, there is (almost) no
undefined behavior.

aka

Low-level does not lie.

aka

This is why we'll focus on it !

Language oddities zoo

```
1 public class Main {  
2  
3     public static void main(String[] args) {  
4         int a1 = 1000, a2 = 1000;  
5         System.out.println(a1 == a2);  
6         Integer b1 = 1000, b2 = 1000;  
7         System.out.println(b1 == b2);  
8         Integer c1 = 100, c2 = 100;  
9         System.out.println(c1 == c2);  
10    }  
11 }
```

OCaml (< 4.05.0)

```
1 | open Nums
2 |
3 | let x = Big_int.big_int_of_int 1 ;;
4 |
5 | x = x ;;
```

```
1 let s = string_of_bool true ;;
2
3 s.[0] <- 'f' ;;
4 s.[1] <- 'a' ;;
5 s.[3] <- 'x' ;;
6
7 1 = 1;;
8
9 Printf.printf "1 = 1 est %b\n" (1 = 1) ;;
```

OCaml

```
1  (##warnings "-3";; (* :-*) *)
2
3  let f x =
4    match x with
5      | true  -> "T"
6      | false -> "F"
7    ;;
8
9  f true ;;
10 f false ;;
11
12 (f false).[0] <- 'T' ;;
13 (f true).[0] <- 'F' ;;
14
15 f true ;;
16 f false ;;
```



```
1 | l = [ s for s in [1, 2, 3] ]  
2 | print(l)  
3 | print(s)
```

JavaScript: got arithmetic ?

```
1 1 / 0
2
3 NaN == NaN
4
5 9999999999999999
6
7 9999999999999999
8
9 "2" + 1
10
11 "2" - 1
12
13 "2" - - 1
```

JavaScript: ==

```
1 [1] == [1]
2
3 [] == ![]
4
5 [] == true
6
7 ![] == true
8
9 2 == [2]
10
11 0 == '0'
12
13 0 == '0.0'
14
15 '0' == '0.0'
16
17 null == undefined
```

More fun

- <https://www.youtube.com/watch?v=et8xNAc2ic8>
- <https://github.com/denysdovhan/wtfjs>

```
1 $x = "2d8" ;  
2 $y = "2d8" ;  
3  
4 ++$x == $y + 1;  
5  
6 print (++ $x . "\n") ;  
7 print (++ $x . "\n") ;  
8  
9 print($y + 1 . "\n") ;
```

```
1 $h1 = md5 ('QNKCDZO') ;  
2 $h2 = md5 ('240610708') ;  
3 $h3 = md5 ('A169818202') ;  
4 $h4 = md5 ('aaaaaaaaaaaaumdozb') ;  
5 $h5 = sha1('badthingsrealmlavznik') ;
```

Which ones are == to each other ?

- A. none
- B. h3 and h5
- C. h1, h3 and h4
- D. *La réponse D*

Scala (until 2.12)

```
1 | List("a", "b", "c").toSet() + "d"
```

Questions ?



<https://rbonichon.github.io/teaching/2021/asi36/>