

Cross-platform compilers with **GNU Lightning**



Paul Cercueil

FOSDEM²⁵

<https://gnu.org/software/lightning/>



« **GNU lightning is a library that generates assembly language code at run-time** ».

<https://gnu.org/software/lightning/>

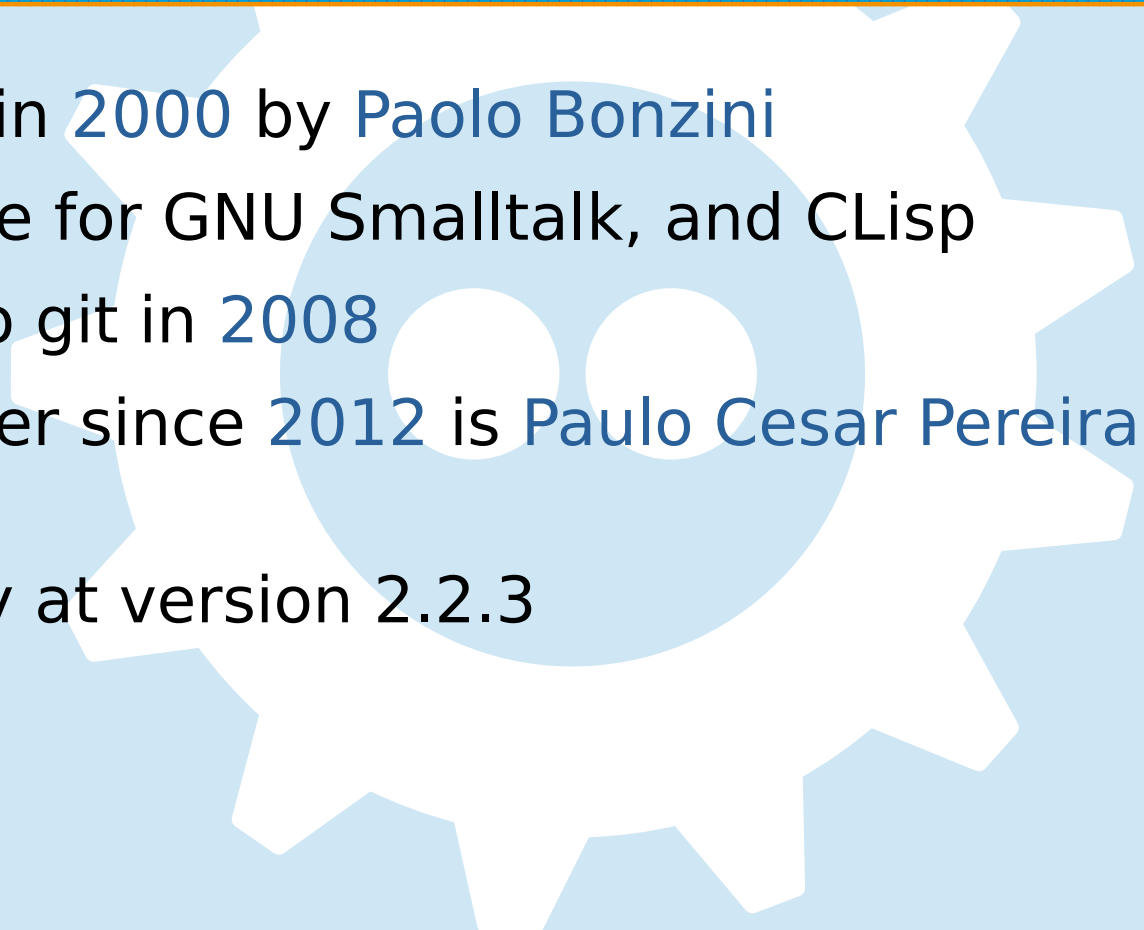


« GNU lightning is a library that generates ~~assembly language~~ code at run-time ».
machine

What is GNU Lightning?

- Code generator part of a JIT engine
- Virtual CPU, with virtual registers and virtual opcodes
- Cross-platform:
 - ARM, Alpha, HPPA, IA64, Loongarch, MIPS, PowerPC, RISC-V, s390, SPARC, SuperH, x86
 - 32-bit and 64-bit
 - Little-endian and big-endian
- LGPLv3, GNU C, no dependencies

What is GNU Lightning?

- Created in 2000 by Paolo Bonzini
 - JIT engine for GNU Smalltalk, and CLisp
 - Moved to git in 2008
 - Maintainer since 2012 is Paulo Cesar Pereira de Andrade
 - Currently at version 2.2.3
- 

Lightrec

- MIPS-to-everything JIT engine for Playstation 1 emulators
- Uses GNU Lightning
- `<ad>`See my talk tomorrow:
 « Writing a dynarec, step by step »`</ad>`

Why Lightning?

- LLVM / libgccjit don't box in the same weight class
 - They were designed for **languages**
 - High-level concepts (e.g. variables)
 - Language-focused algorithms
 - They were designed for **ahead-of-time**
 - Powerful optimizations
 - Generate very fast code... slowly
 - Lightning generates unoptimized code... very fast
- Different tools for different applications.

API overview

- *Minimum* of 6 GP + 6 FP registers:
 - 3 caller-saved: JIT_R0 → JIT_R2 (up to JIT_R(JIT_R_NUM))
 - 3 callee-saved: JIT_V0 → JIT_V2 (up to JIT_V(JIT_V_NUM))
 - 6 floating-point: JIT_F0 → JIT_F5 (up to JIT_F(JIT_F_NUM))
- Simple register allocator (that you don't need):
 - ```
u8 reg = jit_get_reg(jit_class_gpr);
jit_unget_reg(reg);
```



# API overview

- « virtual opcodes » for binary operations, arithmetic, boolean, branching, memory I/O, floating-point math, function calls and function prolog/epilog
- Each instruction is composed of:
  - An operation, like « mul » or « sub »
  - Most times, a register/immediate flag (r or i)
  - A type identifier, when applicable.  
No type suffix = pointer-sized

# API overview

`jit_<op><r|i>[_<type>] (01, 02, 03)`

|                    |                       |                                                 |                    |                       |                                                                    |
|--------------------|-----------------------|-------------------------------------------------|--------------------|-----------------------|--------------------------------------------------------------------|
| <code>addr</code>  | <code>_f _d</code>    | <code>01 = 02 + 03</code>                       | <code>divi</code>  | <code>_u _f _d</code> | <code>01 = 02 / 03</code>                                          |
| <code>addi</code>  | <code>_f _d</code>    | <code>01 = 02 + 03</code>                       | <code>remr</code>  | <code>_u _d</code>    | <code>01 = 02 % 03</code>                                          |
| <code>addrx</code> |                       | <code>01 = 02 + (03 + carry)</code>             | <code>remi</code>  | <code>_u</code>       | <code>01 = 02 % 03</code>                                          |
| <code>addxi</code> |                       | <code>01 = 02 + (03 + carry)</code>             | <code>andr</code>  |                       | <code>01 = 02 &amp; 03</code>                                      |
| <code>addcr</code> |                       | <code>01 = 02 + 03, set carry</code>            | <code>andi</code>  |                       | <code>01 = 02 &amp; 03</code>                                      |
| <code>addci</code> |                       | <code>01 = 02 + 03, set carry</code>            | <code>orr</code>   |                       | <code>01 = 02   03</code>                                          |
| <code>subr</code>  | <code>_f _d</code>    | <code>01 = 02 - 03</code>                       | <code>ori</code>   |                       | <code>01 = 02   03</code>                                          |
| <code>subi</code>  | <code>_f _d</code>    | <code>01 = 02 - 03</code>                       | <code>xorr</code>  |                       | <code>01 = 02 ^ 03</code>                                          |
| <code>subrx</code> |                       | <code>01 = 02 - (03 + carry)</code>             | <code>xori</code>  |                       | <code>01 = 02 ^ 03</code>                                          |
| <code>subxi</code> |                       | <code>01 = 02 - (03 + carry)</code>             | <code>lshr</code>  |                       | <code>01 = 02 &lt;&lt; 03</code>                                   |
| <code>subcr</code> |                       | <code>01 = 02 - 03, set carry</code>            | <code>lshi</code>  |                       | <code>01 = 02 &lt;&lt; 03</code>                                   |
| <code>subci</code> |                       | <code>01 = 02 - 03, set carry</code>            | <code>rshr</code>  | <code>_u</code>       | <code>01 = 02 &gt;&gt; 03(1)</code>                                |
| <code>rsbr</code>  | <code>_f _d</code>    | <code>01 = 03 - 01</code>                       | <code>rshi</code>  | <code>_u</code>       | <code>01 = 02 &gt;&gt; 03(2)</code>                                |
| <code>rsbi</code>  | <code>_f _d</code>    | <code>01 = 03 - 01</code>                       | <code>lrotr</code> |                       | <code>01 = (02 &lt;&lt; 03)   (03 &gt;&gt; (WORDSIZE - 03))</code> |
| <code>mulr</code>  | <code>_f _d</code>    | <code>01 = 02 * 03</code>                       | <code>lroti</code> |                       | <code>01 = (02 &lt;&lt; 03)   (03 &gt;&gt; (WORDSIZE - 03))</code> |
| <code>muli</code>  | <code>_f _d</code>    | <code>01 = 02 * 03</code>                       | <code>rrotr</code> |                       | <code>01 = (02 &gt;&gt; 03)   (03 &lt;&lt; (WORDSIZE - 03))</code> |
| <code>hmulr</code> | <code>_u</code>       | <code>01 = ((02 * 03) &gt;&gt; WORDSIZE)</code> | <code>rroti</code> |                       | <code>01 = (02 &gt;&gt; 03)   (03 &lt;&lt; (WORDSIZE - 03))</code> |
| <code>hmuli</code> | <code>_u</code>       | <code>01 = ((02 * 03) &gt;&gt; WORDSIZE)</code> | <code>movzr</code> |                       | <code>01 = 03 ? 01 : 02</code>                                     |
| <code>divr</code>  | <code>_u _f _d</code> | <code>01 = 02 / 03</code>                       | <code>movnr</code> |                       | <code>01 = 03 ? 02 : 01</code>                                     |

# API overview

## Forward branch

```
jit_node_t *n;

n = jit_beqi(JIT_R0, 0);

...

jit_patch(n);
```

## Backwards branch

```
jit_node_t *n1, *n2;

n1 = jit_label();

...

n2 = jit_beqi(JIT_R0, 0);

jit_patch_at(n2, n1);
```

# API overview

- **Function prolog / epilog:**

```
jit_prolog()
jit_epilog()
```

- **Code generation:**

```
void (*fn)() = jit_emit();
```

- **Code disassembly:**

```
jit_disassemble()
```

# Example of generated code

```
« jit_addr(JIT_R0, JIT_R1, JIT_R2); »
```

```
x86_64:
```

```
 lea (%r10,%r11,1),%rax
```

```
MIPS64:
```

```
 daddu v0,v1,t4
```

```
PowerPC32:
```

```
 add r28,r29,r30
```

```
SH4:
```

```
 mov r2,r1
```

```
 add r3,r1
```

# Demo



**« Unfortunately, no one can be told  
what GNU Lightning is.  
You have to see it for yourself. »**

# Demo

- Imagine a scripting language called « MindBlown »
  - 32K memory cells, each signed 32-bit
  - < / > : switches to the left/right cell
  - + / - : increment/decrement value at current cell
  - [ / ] : repeat until the current cell is zero
  - . : Print the character of the current cell